Infosys® | CONSULTING

NOV 2022 • ISSUE 1

AI & AUTOMATION

BUSINESS JOURNAL

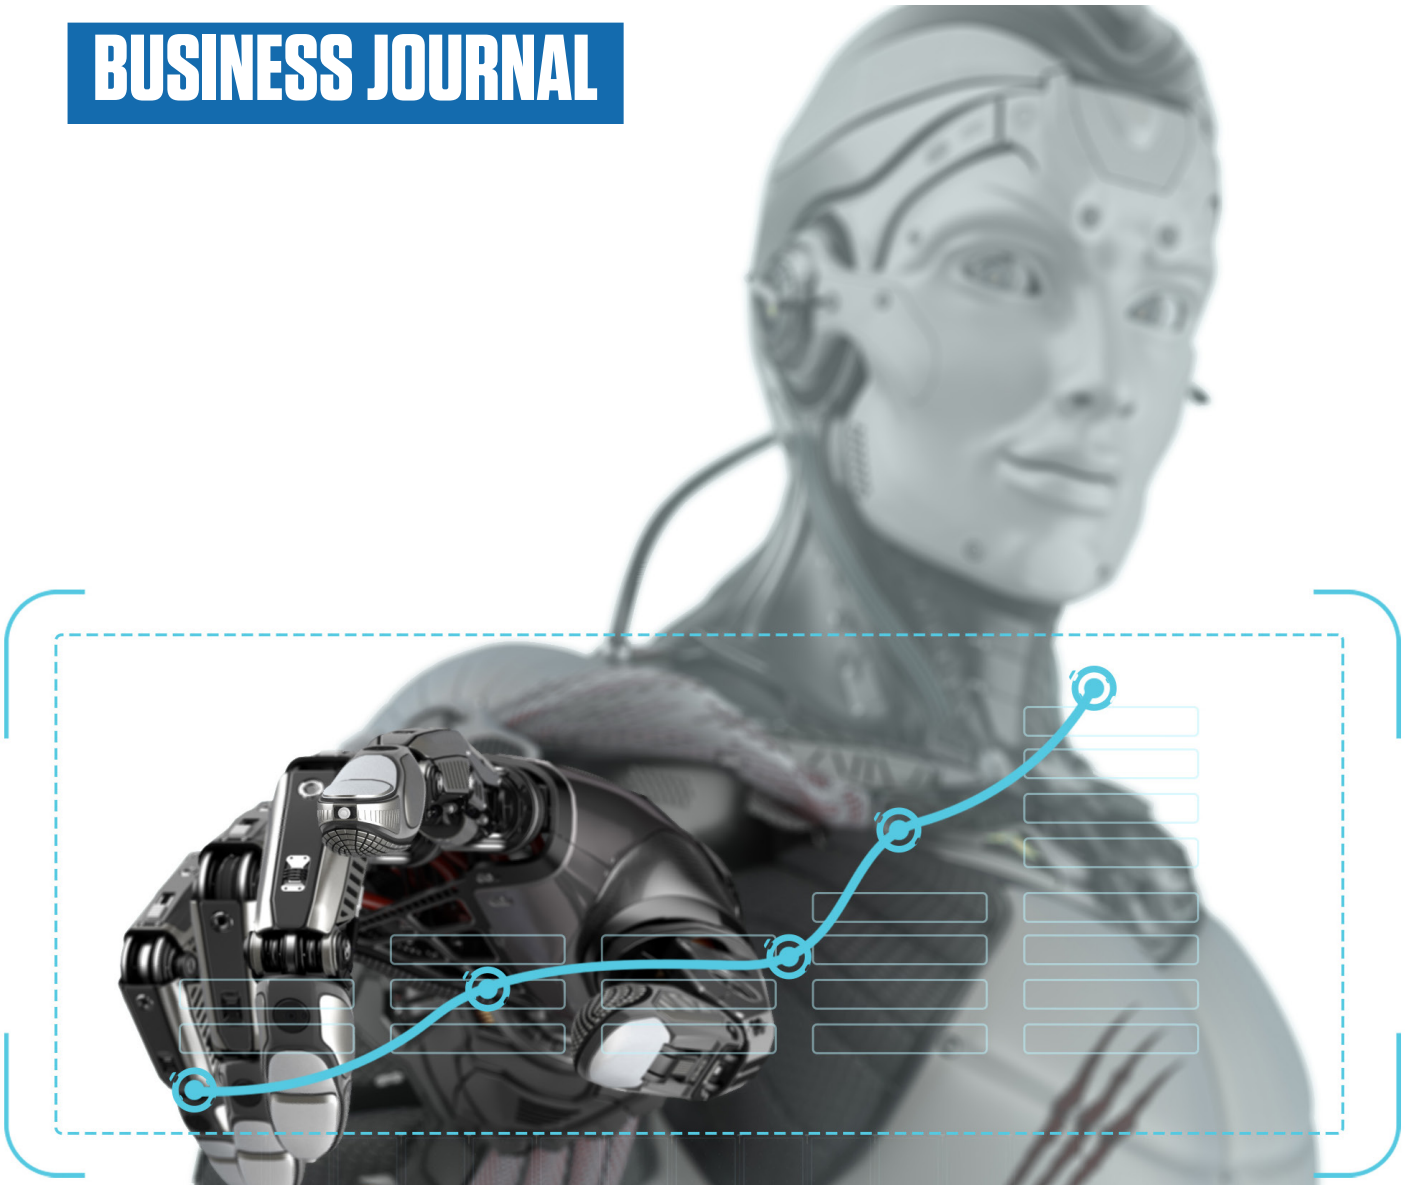

IN THIS ISSUE

01// A NEW FRONTIER IN ANALYTICS: GRAPH DATABASES

02// SCALING AI FOR UNTAPPED BUSINESS VALUE

03// ML FOR BETTER PRODUCT/MARKET SYNERGIES

04// FORECAST ACCURACY FOR BUSINESS VALUE

DEAR AI ENTHUSIAST

At Infosys Consulting, we possess an unwavering belief in the power of technology to unleash new and amazing possibilities to evolve business, amplify human capabilities, and positively impact society.

Welcome to the inaugural edition of the Infosys Consulting's Artificial Intelligence and Automation (AI&A) Business Journal.

As AI and disruptive technology continue to progress, this biannual journal will strive to deliver lessons learned from the academic community, early adopters, industry analysts, and our own experience delivering real-world solutions alongside our clients.

This journal is intended to help inspire and enlighten you on a panacea of future possibilities in this exciting space. It blends theory and practice, drawing a direct line of sight from strategy to implementation – the latter of which is ultimately the barometer of success for any great plan to be operationalized.

Our articles are designed for both the AI strategist and the AI practitioner, with detailed insights on both the promise and the pitfalls of new theories, approaches, and techniques across the spectrum of AI.

We also intend to share lessons learned from real-world experiences, especially around the latest trending topics. We will strive to cite the boundaries of the art of the possible and the realm of the probable. We will share hard fought lessons from the field to sift through fact from fiction. Our team hopes to orienteer you down the path of pragmatic reality so you can achieve the absolute best for your organization.

We hope you enjoy the insights here. I look forward to more stimulating dialogue with you around how AI can help us collectively set a new frontier ahead.

ANDREW DUNCAN
CEO & MANAGING PARTNER
INFOSYS CONSULTING

About the Infosys Consulting AI & Automation Business Journal

As the Artificial Intelligence, Analytics, and Cyber Security Managing Partner, I am excited to share the articles in this inaugural issue of our AI & Automation Journal. We plan to publish the journal biannually where you will find insights, research, and best practices from our work with world class companies we serve as trusted advisors.

Our clients run some of the largest data centers on the planet and are continually pushing the envelope of what is possible for faster, more efficient, advanced AI and machine learning applications on ever-growing mammoth data sets. We enable them to push those limits in environments that ensure the highest levels of security and privacy that their demanding businesses require. In this inaugural issue of the AI & Automation Journal, you will find:

- How to be prepared for graph databases – the fastest growing segment in storage.
- Our approach and process to scaling AI in financial sectors.
- A method using machine learning to reveal the customer mission and innovation opportunities.
- A better way for data scientists to evaluate forecasts for consumer goods.

We look forward to sharing our knowledge with you so that we may remain curious and continue to assist our clients to push the boundaries of what possible is.

SETH SIEGEL
MANAGING PARTNER, AI, ANALYTICS & CYBER SECURITY
INFOSYS CONSULTING

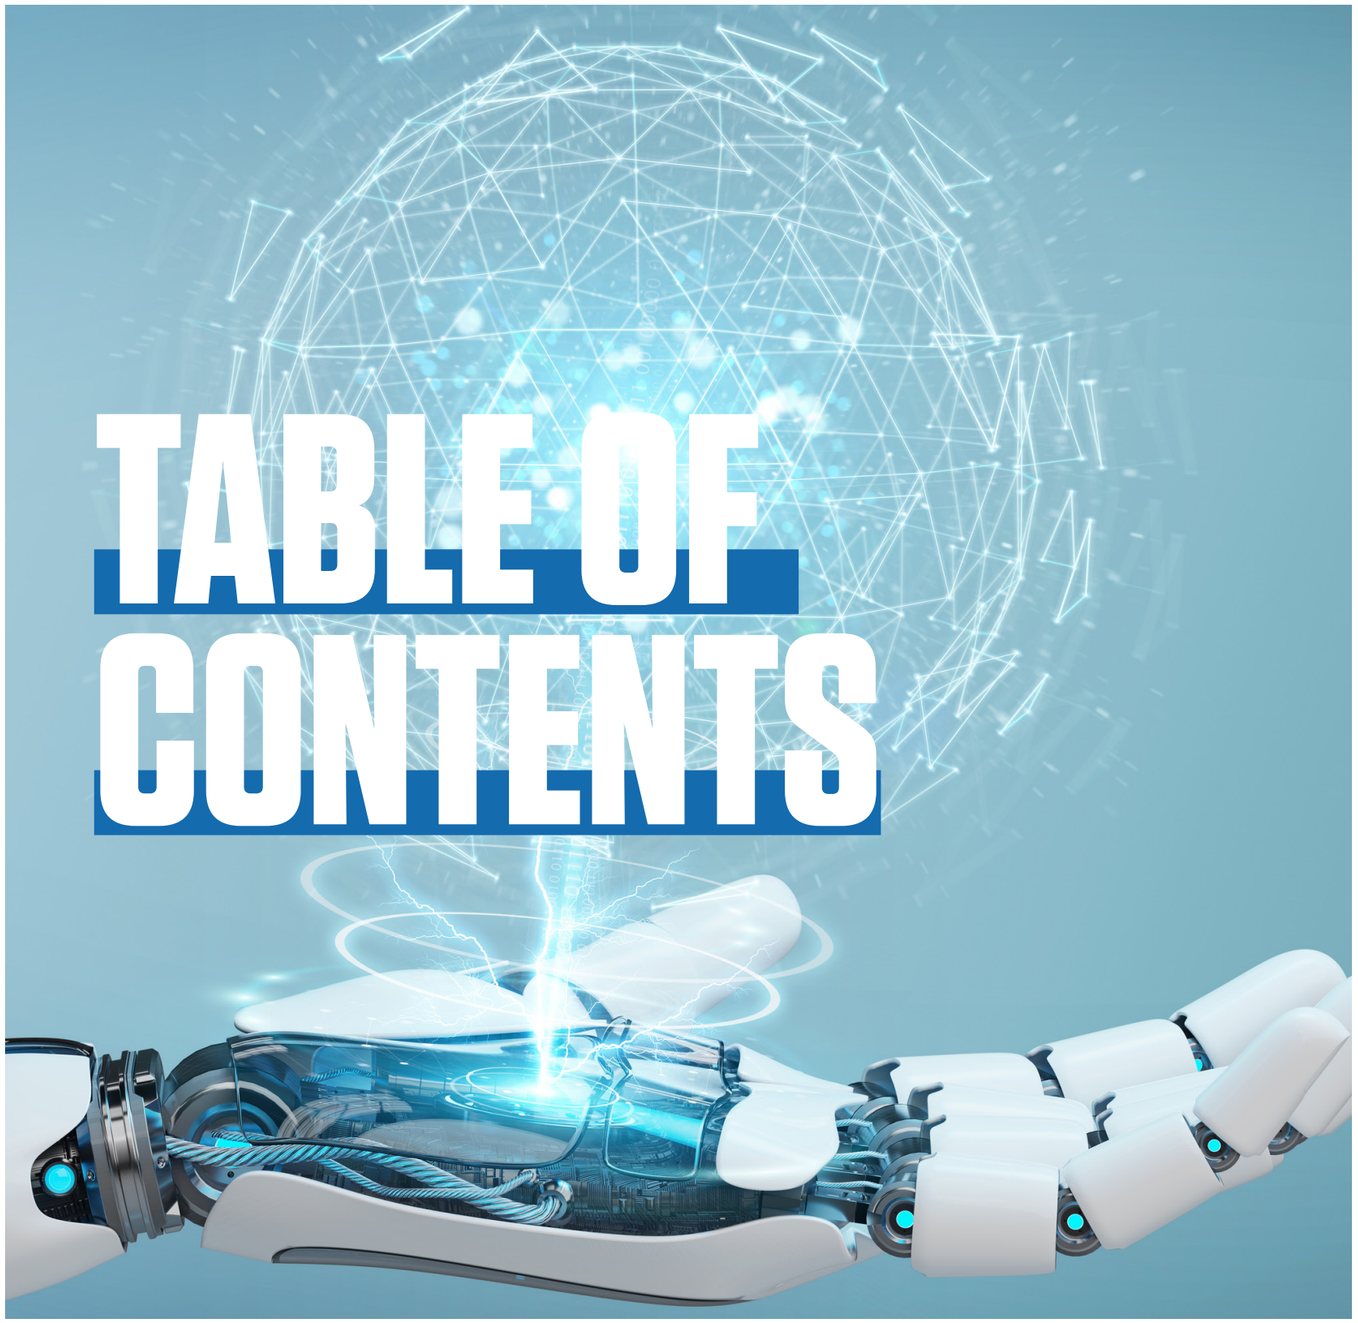

TABLE OF CONTENTS

06// GRAPH DATABASES
DISCOVER A NEW FRONTIER IN ANALYTICS

22// ARTIFICIAL INTELLIGENCE
DOUBLE DOWN ON SCALING AI FOR BUSINESS VALUE

37// MACHINE LEARNING
USING ML TO IDENTIFY PRODUCT/MARKET SYNERGIES

52// FORECASTING
A NEW APPROACH TO FORECAST ACCURACY

GRAPH DATABASES

DISCOVER A NEW FRONTIER IN ANALYTICS

Helping Your Business Stakeholders Change the Game in Front Line Decision Making

BONNIE HOLUB, Ph.D. • VIKRAM RAO • LAUREL SOUZA

Graph databases are an emerging trend in the analytics world. In this article, we outline why they are critical in shaping the next frontier of data science, AI, and machine learning.

A Rapidly Emerging Market Sector will Reshape Enterprise Decision-Making

Companies seeking to monetize their vast data stores quickly realize the true value of their database is not found merely in analysis of discrete data points but rather in analysis of the relationships between the vast array of compiled data. For example, isolating a single fraudster in a database can be the tip of the iceberg in identifying fraud rings for a bank. In 2022, businesses are operating in a more connected world than ever. It is critical to understand the complex relationships and dependencies across various data points to make the best decisions for the enterprise.

A new form of database has been developed to optimize against the increasing complexity of relationships amongst ever-more vast collections of data: the graph database. These specialized databases offer extremely powerful modeling and analytic capabilities that allow businesses to efficiently answer some of the toughest questions – often questions that were previously very hard to address.

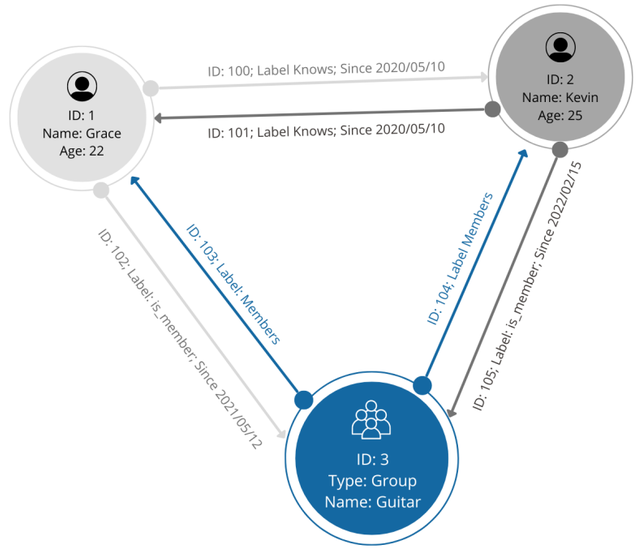

Figure 1: Visual representation of a graph database.

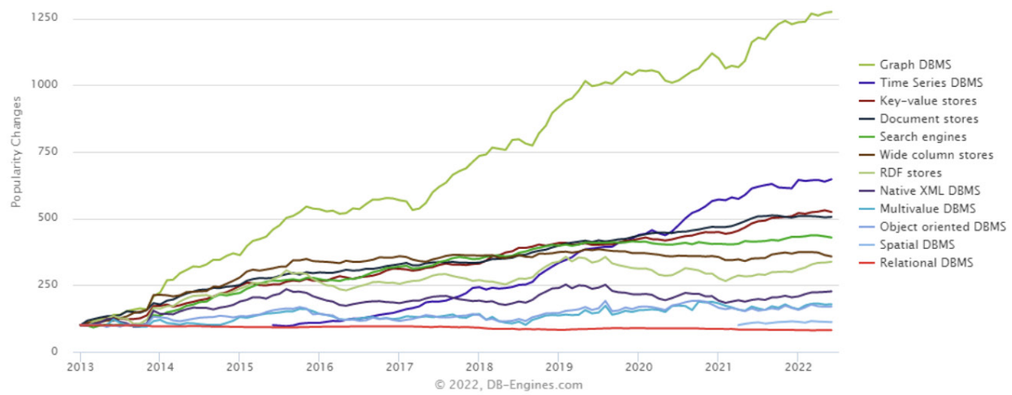

Figure 2: DB Adoption Level Trends (Db-Engines.com, DB-Engines Ranking per database model category, db-engines.com/en/ranking_categories, accessed 2022/06/06).

A 2021 Gartner report identified graph databases (DBs) as a fast-growing market sector, with 80% of data and analytics innovations coming from graph DBs by 2025.[1] Many others have been publishing insights for graph DB adoption levels[2] with the most compelling graphic used to demonstrate the importance of this trend coming from DB engines.com, see figure 2 (above).

This new type of database design is the hidden power behind many of today's most popular applications in daily use: Google search, LinkedIn's degrees of separation, Amazon's "people who bought this also bought" recommendations, email auto-complete, and credit card fraud monitoring, to name a few.

A Powerful New Way to Use Graphs

This is a specific topic in mathematical theory, and specialized databases are being developed to optimize performance for analyzing relationships.[3]

Familiarity with pie charts, line graphs, histograms, and other analytical charts abounds; however, in this article, the term graph is a specific mathematical representation of a network – sets of edges (relationships) and vertices (data points) – describing relationships between data points (see Figure 1, page 7).

A graph database, in short, is a database that is specifically engineered and optimized to store relationships in the mathematical graph form. Relationships across data points are thus valued as equally as the data points themselves, and, as such, these relationships are as accessible as the data itself. Which, in turn, means that relationships can be as easily retrieved and accessed as the data, unlike in a traditional relational database management system (RDBMS), where relationships are stored implicitly.

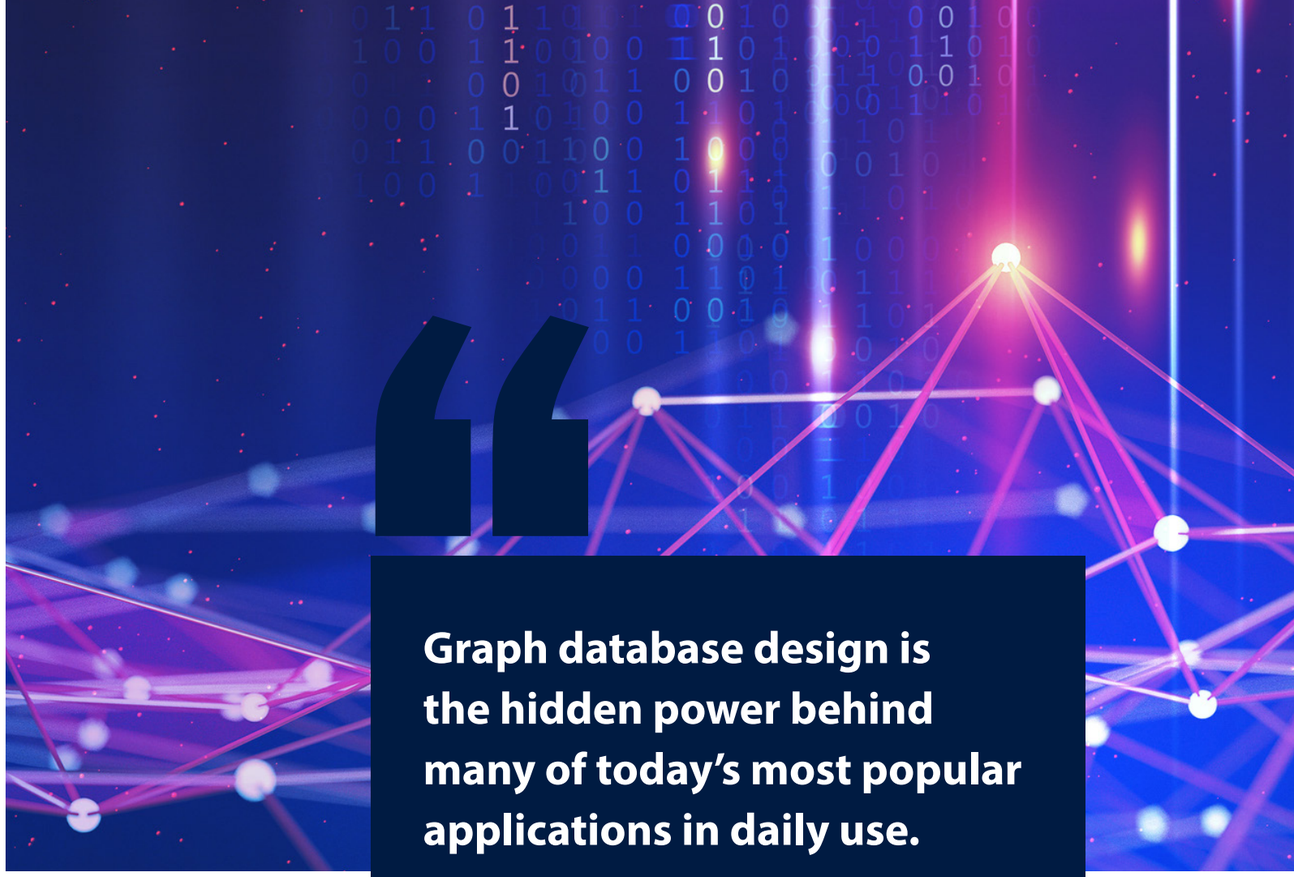

Graph database design is the hidden power behind many of today's most popular applications in daily use.

In addition to natural relationship querying, explicit identification of relationships in the database offers several other advantages: increased retrieval speed, scalability, improved flexibility, and greater agility.

1: Performance

Graph databases provide significant performance improvements when running relationship-driven queries – up to the order of several magnitudes. As more data is introduced and the relationships become deeper in a traditional RDBMS, queries become both more complex and, subsequently, significantly slower. With the relationships between data points incorporated into the structure of the graph database, graph databases are naturally resistant to the performance creep of traditional databases.

2: Scalability

One critical demand of new-age analytics systems is a highly distributed and always-available database architecture. While a traditional RDBMS prioritizes consistency, new-age NoSQL systems, including graph databases, provide an

equally consistent model but concentrate on high availability and distributed systems. This helps organizations to scale quickly and concentrate on high-value implementation.

3: Flexibility
Lacking the rigid structure of an RDBMS, graph DBs are naturally more adaptable to use cases and models that change over time. A graph representation supports continual modification of the database architecture without requiring the costly schema changes required in a traditional RDMBS – when a new property is necessary for a relationship, this property can simply be added to new or existing vertices.

4: Agility
Since graph databases are flexible, it allows the business to be agile. Developers can easily adapt to the ever-changing business requirements without bringing systems down, easily evolving applications. This allows for companies and developers to work in a true agile software development fashion.

Why Should Businesses Care about Graph Databases?

Relationships matter, and analytics derived from them are the next horizon in monetizing data.

Appreciation for the natural reliance of application, use-case, and business scenarios on relationships between data has led leading firms to integrate these relationships into their analyses. Humanity has never been as connected as it is today, and winning businesses harness these connections for the fullest of their potential – often leading to a real competitive differentiator in the market.

So, what can relationships tell us? A lot!
Some typical questions we can answer are:

- How are individuals connected? Are they similar?
- Who is likely to form groups?
- What people are likely to purchase certain products?
- What associations of people are connected to fraud?
- How do parts of an organization interact with each other?

You will notice that majority of data in the modern age exists in the form of relationships between data points with the true business value being the relationship rather than the actual data points. Graph databases are designed with the idea in mind that relationships are just as important as the data itself.

This allows a natural and simplified querying process – unlike SQL queries which typically need very complex and convoluted joins to answer similar questions and thus require some type of analytical or programming expert behind the scenes.

Graph databases are designed with relationship-driven questions in mind, allowing a natural and simplified querying process.

Challenges of Analyzing Relationships with Traditional Relational Databases vs. Graph Databases

Let us explore issues with both SQL and NoSQL databases and how we can overcome the pitfalls with a graph database.

Relational Databases

SQL Databases

- Each business entity is stored on a separate table (e.g., customers, products, item locations, etc.)
- To analyze relationships across business entities, we must create complex table joins:
    - Queries may take a long time to compute
    - Computationally expensive
    - NoSQL databases (key-value databases)

NoSQL Databases

- All data stored on single table
- Relationship queries requiring scanning large number of rows making deep analysis complex

Graph Databases

- Built with relationships in mind
- Entities are naturally connected
- Do not require complex and computationally expensive table joins
- Do not require multiple scans across a single large table
- Mapping relational databases to graph databases

Primary Differences Between Relational Databases & Graph Databases

	Relational Databases	Graph Databases
Organizing Principle	**Tables** of homogeneous data.	**Graphs** of heterogeneous data.
Individual Storage Unit	**Uniform rows** Where each row represents the same types of data, repetitively (e.g., name, address, phone number, contract number for individual after individual).	**Nodes of varying type** A person is linked to another person via a "friend" link, and roommates would each share a "residence" link to the same address.
Cross-individual Properties	**Columns** Data types are stored in columns – address is one column and name is another column.	**Links** Properties are stored in links (called edges) between nodes which have values.
Primary Access Mechanism	**Joins** Bringing together a value from one table to another; a process optimized for batch data processing.	**Traversals** Moving from one node to another over possibly disparate relationships, a process optimized for expressive and flexible exploration of complex relationships.

Complex joins of multiple tables are recognized as challenging for traditional table-based systems which are optimized for repetitive data processing – for example, printing the invoices for each customer in order. An example of the reduced complexity can be demonstrated in the approach to answering the question: "*Which of my contracts is at risk because a competitor has a co-located overseas location?*"

In a traditional SQL database, answering this question would require complex joins of multiple tables – including the accompanying determinations:

- Table 1: Employees (do I have authentication?)
- Table 2: Contracts (which are my contracts?)
- Table 3: Clients (which clients are the contracts with?)
- Table 4: Competitors (who are my competitors on each contract?)
- Table 5: Company addresses (where are the overseas locations of my clients? And where are the overseas locations of my competitors? And where are they collocated?

Conversely, traversals of a graph database can use explicit relationships to directly answer this question.

Let Us Explore Some Compelling Use Cases of Graph Neural Nets

Graph Neural Nets (GNNs) show promise for creating business value in retail, transportation and logistics, and banking.

GNNs have seen advances over the past decade as this efficient approach to supervised machine learning has been applied to quantum chemistry for molecular structures for drug discovery.[4] This work has been extended to larger data sets covering a variety of applications, including the three major sectors of retail, transportation, and logistics,[5] and banking and financial services.

Retail Sector

Product recommendation is one of the fields in retail where GNNs can be leveraged. Compared to traditional methods, which rely on regimented tables of row/column data marching through a data processing system, GNNs can use both the product information (user or item based) in the heterogeneous graph database. This more holistic view of the data and relationships dramatized the interactions between the users and products.

A user-product interaction is learned by the GNN applying information from the nearby nodes, and the GNN-derived model determines the best product to be recommended to a particular user.[6] The advantage of using GNN for product recommendation is it can learn from raw data without manual intervention creating feature significance *a-priori*.

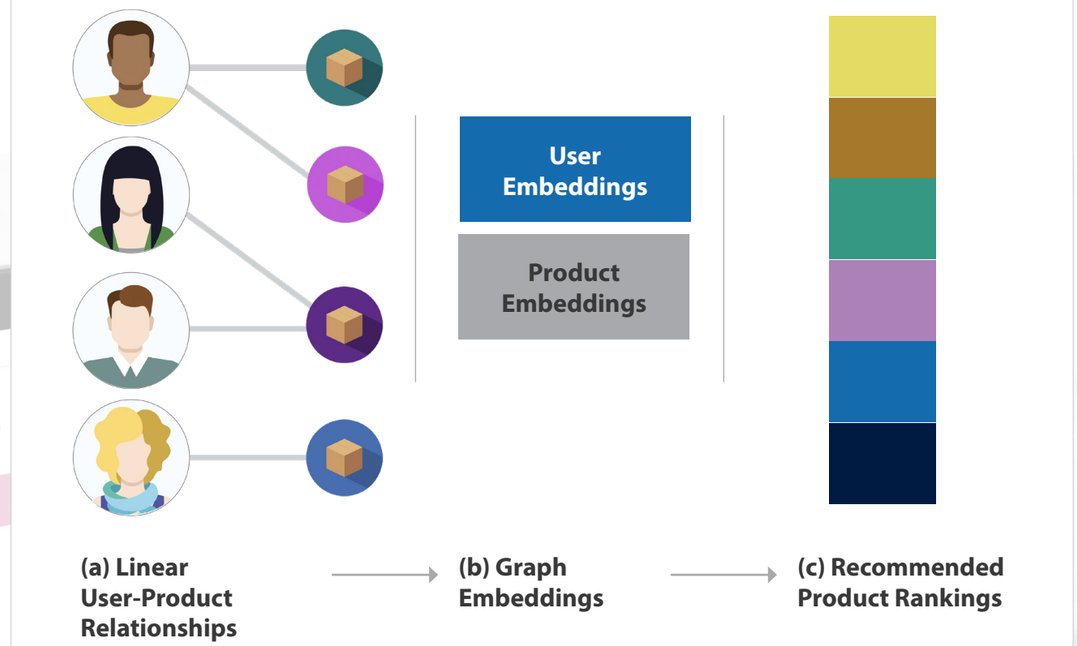

Figure 3: Based on the product purchases of users, a knowledge graph is created (a), graph embeddings consisting of the user-product relationships are performed (b) during model building. The GNN derived model ranks and outputs the recommended product to the users (c).

Figure 4: A sample road network is partitioned into segments (a), each segment is considered a node (b) and the corresponding nodes are connected by edges and form a super-segment graph, (c) GNNs are trained on these super segments to find optional routings, construction and traffic workarounds, and other logistics problem solutions.

Transportation & Logistics

Estimated time of arrival (ETA) is an aspect that is very useful in the transportation industry, especially in air travel and land transportation services. The main challenges in this field before the advent of the GNN, were the shortcomings of computational power and the limited amount of real-time data to solve these computationally intensive problems. Recent advances are extending GNNs to this domain with promise.[7]

GNNs are presently being used in the ETA prediction for land transport, by taking both the historical and real-time data into consideration. The roadmap is segmented and the flow of traffic within these segments are predicted, thus forming a chain of predicted segments representable in a graph database. Each road segment can be a few meters in length, and they are enclosed within larger super-segments.[8] Factors such as the road width and conditions, traffic conditions, historical and present data, along with detours are taken into consideration within each segment.

GNNs are better able to identify traffic congestions than the traditional methods by utilizing the traffic information from the super-segments. Multiple super-segments can span any number of road segments to assess various common routes that pass through that road segment. Google Maps makes use of GNNs and frequently serves huge amount of travel time queries.[9]

Banking & Financial Services

In this sector, determining potential customers from the existing ones helps banks to increase revenue and explore more opportunities.[10] Better understanding of the customer network creates an opportunity for the bank to provide new offerings like cash management and liquidity solutions to likely potential customers. It also helps in customer loyalty, lower default rates and fraud prevention.[11]

GNNs can assess the existing graph of bank customers and determine the relationships and similarities between these customers. The new potential customer can have similar characteristics with the bank's existing customers and could also have similar mutual companies (clients) related to each other. In addition, using GNN can lead to an analysis of the bank's inter-bank transactions, and the possible reduction of Days Sales Outstanding (DSO) and the improvement of Accounts Receivable (AR) for the corresponding banks. Grouping the similar networks of correspondent banks facilitate the identification of the most profitable client-supplier networks and the key customers within these networks.

Figure 5: DB Adoption Level Trends (Db-Engines.com, DB-Engines Ranking per database model category, db-engines.com/en/ranking_categories, accessed 2022/06/06).

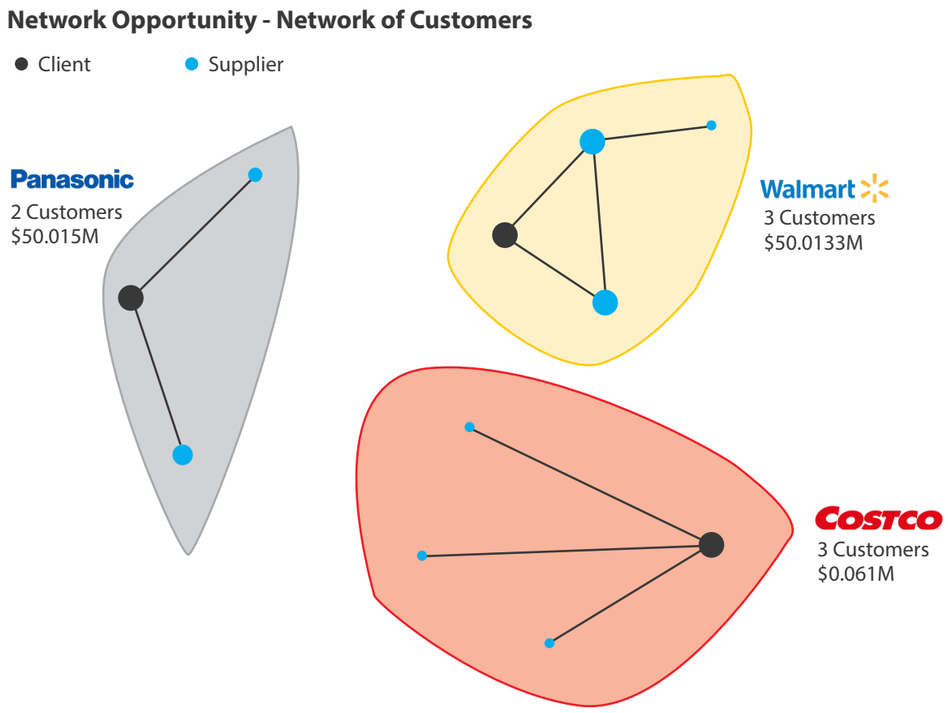

Graph Networks in Operation

LinkedIn

This image represents a graph network of one of the author's links in LinkedIn. It shows groups of individuals affiliated with an institution all in the same color, a large group of densely connected people at the top in dark blue, another, more isolated group in light blue on the right, a cluster of red on the left, and smaller groups of magenta and orange left to right across the bottom. This is the kind of network information LinkedIn uses to suggest new links, "here are some people you may know."

Figure 6: Author's LinkedIn graph network.

Panama Papers

The story of corrupt investments referred to as "The Panama Papers" broke in 2016 because journalists used a graph network to analyze the relationships among investors, officers, and clients of companies incorporated in Panama. Here is an example using then president of Azerbaijan Ilham Aliyev's family.[12] It shows their connections to banking, insurance, construction companies and more.

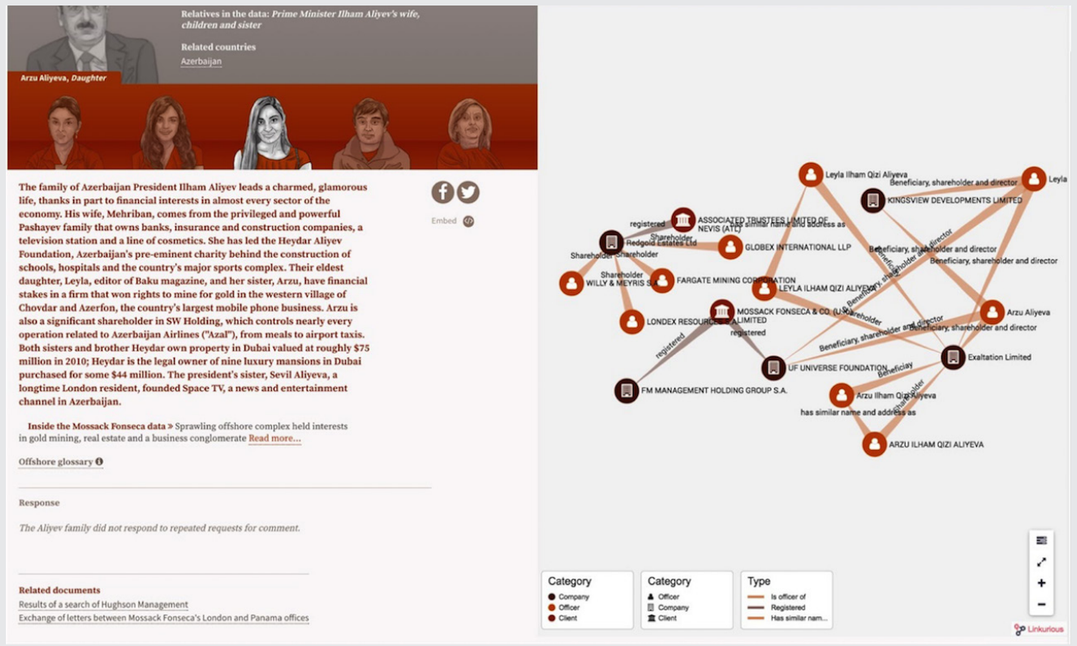
Relatives in the data: Prime Minister Ilham Aliyev's wife, children and sister
Related countries
Azerbaijan
Arzu Aliyeva, Daughter
The family of Azerbaijan President Ilham Aliyev leads a charmed, glamorous life, thanks in part to financial interests in almost every sector of the economy. His wife, Mehriban, comes from the privileged and powerful Pashayev family that owns banks, insurance and construction companies, a television station and a line of cosmetics. She has led the Heydar Aliyev Foundation, Azerbaijan's pre-eminent charity behind the construction of schools, hospitals and the country's major sports complex. Their eldest daughter, Leyla, editor of Baku magazine, and her sister, Arzu, have financial stakes in a firm that won rights to mine for gold in the western village of Chovdar and Azerfon, the country's largest mobile phone business. Arzu is also a significant shareholder in SW Holding, which controls nearly every operation related to Azerbaijan Airlines ("Azal"), from meals to airport taxis. Both sisters and brother Heydar own property in Dubai valued at roughly $75 million in 2010; Heydar is the legal owner of nine luxury mansions in Dubai purchased for some $44 million. The president's sister, Sevil Aliyeva, a longtime London resident, founded Space TV, a news and entertainment channel in Azerbaijan.
Inside the Mossack Fonseca data » Sprawling offshore complex held interests in gold mining, real estate and a business conglomerate Read more...
Offshore glossary
Response
The Aliyev family did not respond to repeated requests for comment.
Related documents
Results of a search of Hughson Management
Exchange of letters between Mossack Fonseca's London and Panama offices
Embed
Leyla Ilham Qizi Aliyeva
Beneficiary, shareholder and director
Leyla
KINGSVIEW DEVELOPMENTS LIMITED
ASSOCIATED TRUSTEES LIMITED OF NEVIS (ATC)
registered
Shareholder
Redgold Estates Ltd
GLOBEX INTERNATIONAL LLP
Shareholder
Beneficiary, shareholder and director
Shareholder
WILLY & MEYRIS S.A.
FARGATE MINING CORPORATION
LEYLA ILHAM QIZI ALIYEVA
MOSSACK FONSECA & CO. (U.K.) LIMITED
registered
LONDEX RESOURCES LIMITED
registered
Arzu Aliyeva
Beneficiary, shareholder and director
Exaltation Limited
LIF UNIVERSE FOUNDATION
FM MANAGEMENT HOLDING GROUP S.A.
Arzu Ilham Qizi Aliyeva
has similar name and address as
ARZU ILHAM QIZI ALIYEVA
Category
Company
Officer
Client
Category
Officer
Company
Client
Type
Is officer of
Registered
Has similar nam...
Linkurious

Figure 7: Graph database of The Panama Papers relationships, "The Panama Papers: Exposing the Rogue Offshore Finance Industry - ICIJ International Consortium of Investigative Journalists, "The Panama Papers: Exposing the Rogue Offshore Industry," accessed 2022/10/14.

Advanced Implementation of Graph Databases

Automating Skill Matching for Recruiting

Figure 8 below shows the relationships between a person/ potential employee, their skills, occupation, and certifications, providing proof of skill acquisition to be matched against job postings for automated skill matching.

Figure 9 to the right shows a graph database representation of an ontology, a network of concepts and their relationships, which can be used to match candidates to job openings.

A graph database can match candidates – whose skills have been automatically extracted from a resume or online portal (like LinkedIn) – with the network of skills required for a job.

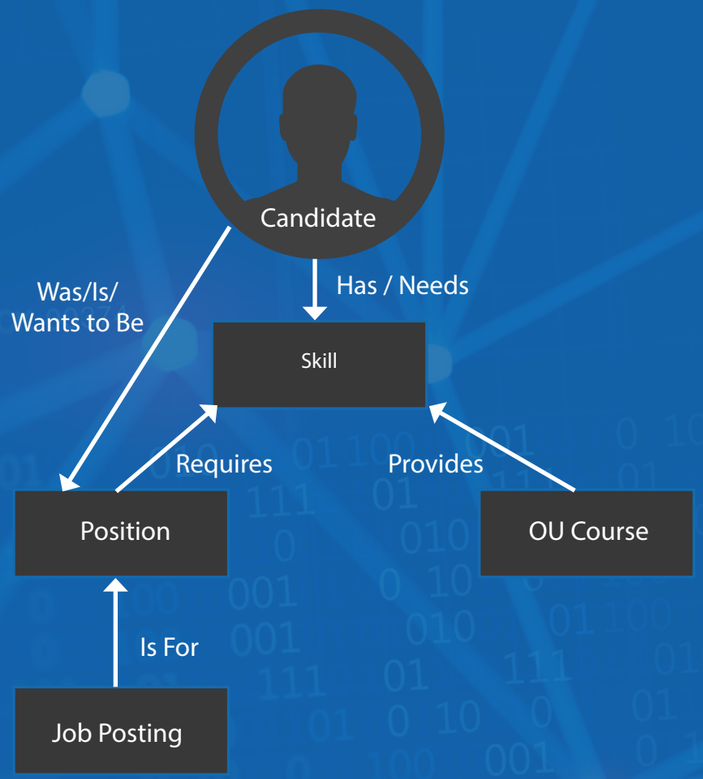

Figure 8: Graph database representing the potential candidate relationships.

Figure 9: Illustrative example of using a graph database to match candidates to job openings.

Areas to Consider in Graph Database Implementation

We have discussed graph databases and their usefulness in the next frontier of analytics, data science, AI, and machine learning. We discussed their applicability to a variety of use cases across several industries and showed examples of implementations. We leave you with a series of points to consider in graph database implementation.

- Will the graph database meet my growing needs as data collection in the enterprise expands? Key offerings in the market include: Neo4j, TigerGraph, Amazon Neptune, Microsoft's Azure Cosmos DB Graph, and Oracle.[13]
- Can the graph database be used in the cloud? On premise? Hybrid?
- Can the graph database work in a distributed computing environment?
- Does it scale, and descale well?
- What is the total cost ownership? This should include the ability of the graph database to:
    - run on multiple servers
    - to scale up and down as processing demands
- It should also include calculations for security including:
    - data lineage, data governance
    - deployment speed and ease
    - as well as associated development tools like data discovery, data preparation, visualization

Graph Databases: A Powerful Enabler for Future Success

Graph data platforms enable efficient analysis of relationships within and among data points in a way that traditional database management systems do not. This relational graph data is recognized as the next significant frontier for monetizing data. And, if your organization can turn its data in a monetization engagement, you will be certain to crack the ear of top management. Best-in-breed tools can simplify deployment, maintain the highest security standards, scale efficiently on-premise, in the cloud, or in hybrid systems enabling thousands of concurrent users and queries per seconds to engage advanced machine learning tools, with visualized results.

About the Authors

BONNIE HOLUB, PH.D.

Associate Partner; Minnesota

Dr. Holub is an Artificial Intelligence and Data Science evangelist and leader who collaborates with diverse constituents and grows businesses by correlation disparate sets of big data for actionable results. She has decades of experience as a consulting professional, tenured faculty member, entrepreneur, and board member.

VIKRAM RAO

Senior Principal, AI & Automation; Bangalore

Vikram has over twenty years of experience in software development, design, product development, consulting, innovation, and R&D. Vikram is passionate about applying new technologies to solve complex business problems. Vikram has over ten years of experience in advanced analytics using machine learning, deep learning, computer vision, graph database and generative modeling.

LAUREL SOUZA

Senior Consultant, AI & Automation; Bangalore

Laurel has worked in AI across machine learning, chatbots and computer vision. He has 10 years of experience in computer science and has delivered various AI projects in retail, energy, financial services, and logistics. He has practical experience in various programming languages. He is certified in deep learning and TensorFlow, and has done his Master of Technology (M.Tech) degree in computer science & engineering.

ARTIFICIAL INTELLIGENCE

DOUBLE DOWN ON SCALING YOUR AI

Capture Immense, Untapped Business Value

RYAN MARTIN

Most organizations have yet to scale the true potential of artificial intelligence. In this article, we introduce vertical AI and why it offers businesses an opportunity to "go big" with AI.

103%

Increase in undergrad courses that teach AI skills, in the last four years.

$93B

Private investment in AI in 2021, double that of 2020.

30x

Increase in AI patents filed between 2015 and 2021.

AI is Ready for a New Frontier of Growth in Enterprise

Recent years have provided several accomplishments in expanding artificial intelligence (AI), and the hype is real when it comes to expanding AI capabilities across almost every business sector. From expansion of the AI technology stack, capabilities in interpreting unstructured data, costs to train models, available training both in higher education and professional resources and expanded funding – the golden age of AI is certainly upon us.[14]

Education is more available than ever. An AI Index survey conducted in 2020 suggests that the world's top universities have increased their investment in AI education over the past four years.[15] The number of courses that teach students the skills necessary to build or deploy a practical AI model on the undergraduate and graduate levels has increased by 103% and 42%, respectively, in the last four academic years.

The funding space confirms this excitement. According to the Stanford AI Index, "private investment in AI in 2021 totaled around $93 billion—more than double the total private investment in 2020."[16] Crunchbase data reveals that 2022 is off to a roaring start, with 559 individual funding rounds for AI-related companies as of July 2022.

AI innovation continues to expand capabilities: the number of AI patents filed in 2021 is more than 30 times higher than in 2015, showing a compound annual growth rate of 77%.[17] "AI innovation is happening at a rapid pace, with an above-average number of technologies on the Hype Cycle reaching mainstream adoption within two to five years," said Shubhangi Vashisth, senior principal research analyst at Gartner.[18]

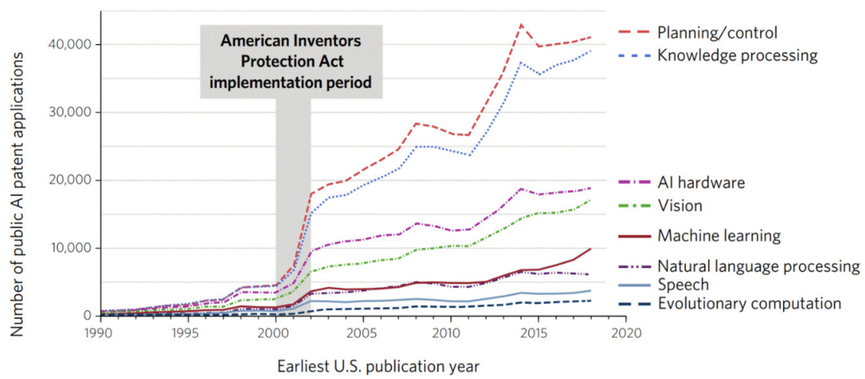

Figure 10: shows the volume of public AI patent applications by year, categorized by AI component in the U.S. The eight AI components in FIG. 10 are defined in an article published in 2020[19] by the USPTO. Most of the AI components have experienced explosive growth in the past decade, especially in the areas of planning/control and knowledge processing (e.g., using big data in automated systems). Also represented is the significant uptick in applications following the implementation of regulatory protection through the American Inventors Protection Act. Source: SheppardMullin.[20]

Most organizations have started adopting AI or are looking to do so. According to a 2021 Pactum survey, 97% of senior executives planned to invest heavily in AI. Additionally, 61% of business professionals cite machine learning and artificial intelligence as their company's most important data-driven initiatives for the next year.[21]

But most organizations are yet to scale to the **true potential** of AI.

Some statistics help paint this clouded picture. In a 2019 MIT Sloan Management Review and Boston Consulting Group AI survey, 7 out of 10 companies reported minimal or no value from their AI investments.[22] That has started to change, with 92% of large companies reporting that they are achieving returns on their data and AI investments, according to a 2022 survey of senior data and technology executives by NewVantage Partners.[23] Note that this figure is limited to large companies, and we find small to mid-sized companies have typically trailed in investment and are earlier in their journey to achieving returns.

If we dig further into the data, we see challenges in moving toward widespread adoption and being a data-driven organization when it comes to AI. There has been a massive uptick in adoption, and 96% of organizations say that they are at least doing *something*, with AI initiatives that are underway in a pilot of limited production. Only 26% report that adoption is widespread, and just 26% identify themselves as a data-driven organization.[24]

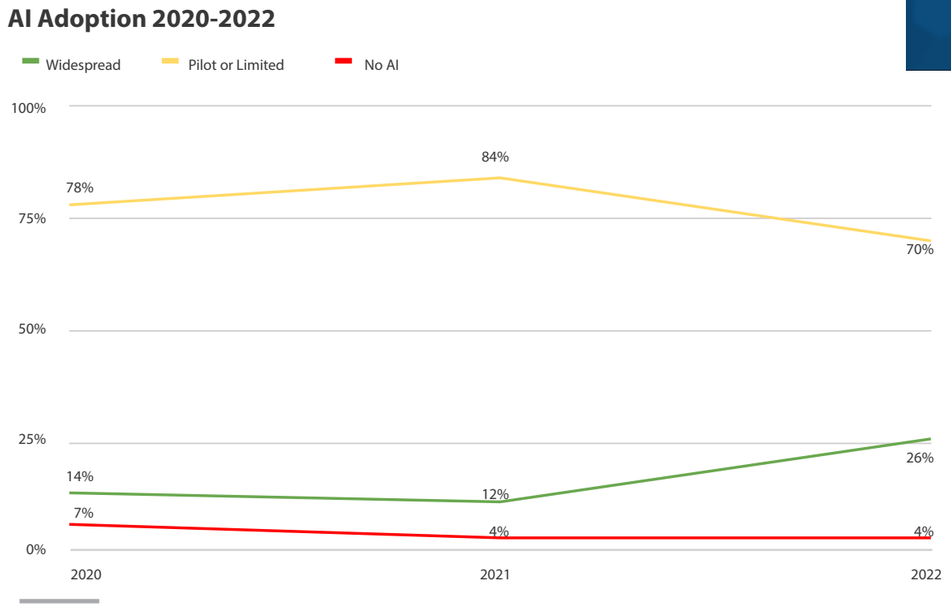

Figure 11: State of AI Investment and Adoption, 2020-22. Source: NewVantage Partners.

92%

of senior data and technology executives at large companies report that they are achieving returns on their data and AI Investments

Source: NewVantage Partners

Moving beyond a pilot is not the ultimate goal. It is a step on the path to achieving widespread adoption and *impactful vertical AI use cases that transform your business and the customer experience*. The opportunity to go big can be addressed by focusing on Vertical AI[25] – taking a holistic view of related use cases, addressing a business problem top-to-bottom. This top-to-bottom vantage point moves beyond solving small problems using broad data science tools (horizontal AI).

In this article, we focus on the opportunity to go truly vertical with AI – for example, applying AI across related, stacked use cases in a customer's lifecycle with your organization. Or the lifecycle of a product from supply chain, through manufacturing, to distribution. We will look to a non-traditional industry for insight on the art of the possible in compounding value through vertical AI use cases. And finally, we will lay out components of an approach to frame this holistic viewpoint of vertical AI, and what organizations need to first consider when going vertical.

Going Vertical

To achieve vertical AI and realize the compounding impact it can have across use cases, we must think more holistically about how we approach, solution, and govern AI. With the lessons from the pandemic, the next years are the opportunity to go big, scaling beyond small "digital transformation." For most established organizations, this means more than simply "thinking bigger," as we have legacy technologies to consider, the established security protocols and compliance requirements, and how to align teams and talent to the best solution.

Below we lay out three key components for taking this holistic approach, which are fundamental to your organization going vertical with AI: strategy, platform, and enablement.

It's our belief that taking a more holistic viewpoint of the potential for AI – a broader viewpoint of the value proposition, a more comprehensive view of the use cases that could achieve it, and technology-second stance on its solutioning – will unleash compounded value across our transformations. More holistic integration of artificial intelligence data and models, across multiple use cases and domains, can greatly compound the value we achieve in applying these technologies. Furthering customer experience, enhancing business insights, and driving more efficient compliance.

And what better place to look than the world's oldest trade – agriculture. We wouldn't traditionally think of farming as thought leaders in the space, but it may paint a picture for this data-driven approach, and the art of the possible for vertical AI. It may be the impetus we all need to look more holistically.

A Brief History of Vertical and Horizontal AI Terminology

Vertical AI and horizontal AI are terms that have been part of use case discussions over the past few years. The concept behind these definitions lies in whether model(s) are built to apply broadly across a particular application (horizontal AI) versus specifically tailored to "stacked" use cases for a top-to-bottom business challenge (vertical AI). In its application, horizontal AI tends to be a common deployment strategy and providing general data science capabilities that are flexible across multiple applications, such as language processing. A good example is Apple's Siri or Amazon's Alexa – they specifically target spoken human language (as of this writing, Alexa covers eight languages) interpretation to answer a nearly limitless set of questions or needs.

Vertical AI is specific and tailored to a business challenge, focusing not on simple, narrow use cases but instead combining or stacking multiple use cases and technologies to achieve a greater value proposition. An example of vertical AI in consumer banking may be the use of chatbots during new customer onboarding to gather information and customer needs. This combines with automated product recommendations and hyper-personalization, running alongside know your customer (KYC) and fraud modeling. This is followed by automated account opening, portfolio recommendations in brokerage relationships (robo-advising), including AI-driven underwriting for loan and insurance products. The ongoing relationship is improved with AI customer insights around behavior, potential for new or personalized products, and an integrated, continuous, data-driven risk management and fraud process. We seek to break down the silos between front and back office, or lines of defense in risk management, to comprehensively manage the relationship while providing hyper-personalization, improved service, and better compliance.

We also see AI terminology described from an infrastructure view, in the sense of providers building *integrated* AI stacks of technologies. The platform capabilities are important to achieve vertical AI, in the sense of having the technologies available to process use cases as described above. This is a component of the "platform" component in the approach we describe within this paper.

The concept we explore here leverages both vertical and horizontal AI – but focused more on the viewpoint we take in their application. We're focusing on a holistic integration of use cases and models, particularly where we see the potential for *compounding* value and ancillary value impacts across the lifecycle of a person or product with which we interact. We call this vertical AI in the truest sense of the definition.

Taking a Lesson in the Art of the Possible, from Farming

The farming industry is believed to be one of the oldest, if not the oldest, profession in the world. And from humble beginnings using hand-made tools, it is also an industry that is recently booming with technological innovation.

We see a breadth of AI applications in specific use cases in agriculture and farming, such as[26]:

- Crop and soil monitoring
- Insect and plant disease detection
- Livestock health monitoring
- Intelligent spraying
- Automatic weeding
- Aerial survey and imaging
- Produce grading and sorting

Taking just one use-case – crop and soil monitoring, for example – soil samples are analyzed chemically to determine soil chemical composition, such as nitrogen dioxide, aluminum, magnesium, and chlorine levels – importantly ammonium, the compound converted to nitrates and nitrites by bacteria. Soil is also analyzed by computer vision for composition analysis, such as sand, clay, and rocky material. Some models also further combine this with weather data, time since fertilization, pH, and soil conductivity measurements. It uses this data to predict how much total nitrogen the soil has now and how much it will have up to 12 days in the future, to predict the optimum time for fertilization.[27]

If we stopped there, we've already generated significant real value to the farmer and beyond. The farmer can deliver the right chemicals, at just the right quantity and the right time. Resulting in less fertilizer usage, which

AI IN AGRICULTURE: SOIL MONITORING

Soil samples are analyzed to determine soil chemical make-up, composition, further combining weather data, time since fertilization, pH levels, and soil conductivity. This information is used to determine timetables for agricultural activity.

makes healthier crops, potential for better yields, and a healthier environment. Over-fertilization has so far rendered 12% of once-arable land worldwide unusable and the use of nitrogen-based fertilizer has risen by 600% in the last 50 years.[28]

But the use case doesn't stop there if we move beyond a targeted use case and into vertical AI. The soil composition and appropriate application of fertilizer is just the beginning. There are numerous other opportunities for deploying AI across the value chain for farmers and, specifically, in the lifecycle of a crop:

Crop analysis: Crop analysis through computer vision technology, either land or drone based. Analyzing for potential disease, pest control issues, ripeness, readiness for harvest.

Soil monitoring: In-ground soil moisture monitors and thermal imaging cameras, interactive with AI-driven watering systems to precisely target water application. Under and over-watering both stress crops, reducing yields and increasing the risk of disease. Runoff makes it much more likely the chemical applications are less effective and more likely to pollute the environment. Fresh

water is also one of the scarcest resources in the world – agriculture is estimated to account for 70% of freshwater usage worldwide.[29]

Automated equipment: Automated equipment, such as drones, harvesters, combines, and tractors can be deployed to apply chemicals, remove weeds, harvest ripe crops, dispose of diseased or bad crop, and transport materials within the farm and beyond.

Automated supply chain: Automated sorting, quality scanning, packaging, shipping.

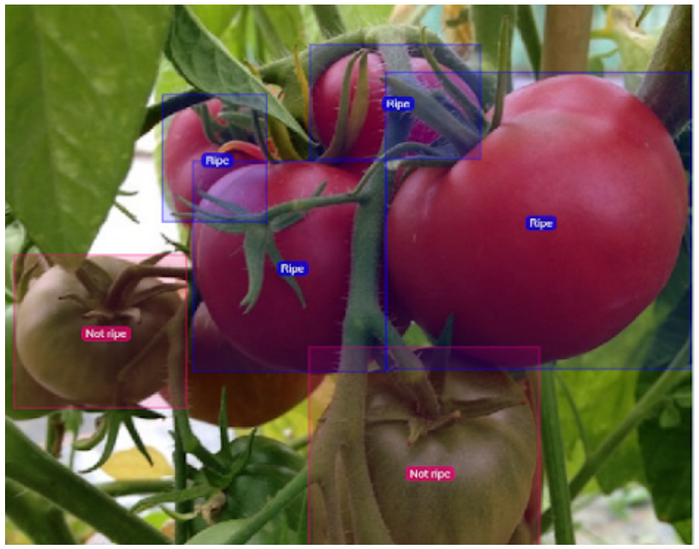

Figure 12 (left): An illustrative weed and mature corn field annotation utilizing drone technology. Figure 13 (right): Illustrative detection of ripe vs. non-ripe tomatoes utilizing computer vision.

Automated soil samples are analyzed and interpreted through AI models to determine irrigation and fertilization needs, as well as best crop rotation to maintain soil quality. Computer imaging interprets pictures of crops to identify pest management concerns, crop health, and optimal timing for harvesting.

At a farmer-level, this results in healthier crops, lower chemical usage, lower water usage, and the potential for better income. On top of that, some results point to significant improvement in crop yields, exceeding 30% improvement for the same area of farmland.[30]

At a consumer level, this results in healthier food and better food choices, likely more availability of fragile or difficult-to-raise crops, and better prices due to better crop yield. At a societal level, we utilize less fresh water, and have less of an impact on the environment due to lower fertilizer and pesticide usage.

92%

of executives point to culture as the greatest impediment to developing a corporate data strategy

The compounded effect of linking use cases together through vertical AI – building through shared data and analysis through the lifecycle of agriculture – is clear. Farmers continue to innovate and automate, with the future holding vertical farming as an ultra-automated solution to overstretched farmland.

But moving beyond farming into other sectors, how can we take the lessons learned to realize the true potential for vertical AI?

Accelerating Toward Vertical AI

Taking the concept of vertical AI across non-agriculture industries has a similar opportunity to compound the value proposition. In service-related industries, we can connect use cases across functions, following a customer through their full lifecycle, from the front office to the back office, and providing a better customer experience. Risk and compliance can occur at the edges, alongside real-time decision making powered by AI, rather than after through second or third lines of defense. In product companies, we can connect across the product lifecycle, from supply chain through manufacturing, quality, and environment monitoring, through packaging, warehousing, and distribution channels.

Getting to vertical AI requires more than just a broader viewpoint of use cases and the ability to connect them across organizational silos. It will require organizations to unleash the potential for this technology through the right holistic strategy, combining it with the technology stack to enable, security and risk posture, talent, and teaming in place to step across silos. Ultimately, the goal is to view the operating model as existing across several platforms, combining technology capabilities with end-to-end business process ownership and the related assets and responsibilities.

Below we lay out the three components across which this holistic viewpoint applies, to keep in mind as you continue this journey to widespread adoption and vertical AI.

1. Strategy - More than half of organizations have developed a corporate data strategy – 53.0%.[31] Over 90% of executives (92% in 2022) point to culture as the greatest impediment to achieving this business outcome.

But those strategies in many cases have yet to fully flush out the benefit of AI and move toward widespread adoption. “One can extrapolate…that most

organizations have yet to develop a corporate data strategy that correlates directly to successful business outcomes and the realization of business value."[32]

A PwC survey shows that organizations taking a more holistic approach are far more likely (36% versus 20%) to report widespread AI adoption than those that took a piecemeal approach. "They're roughly twice as likely to report substantial value from AI initiatives to improve productivity, decision-making, customer experience, product and service innovation, employee experience and more. By bringing so many leaders together from across the organization, a holistic approach facilitates scale and data sharing."[33]

To take a more holistic approach to vertical AI, we recommend thinking through the following considerations:

Zoom out

While staying narrowly focused may help identify a specific use case for AI, it tends to lead us to be blinders-on playing whack-a-mole. This leads to disjointed use cases that don't seem to connect or serve a higher good.

Take a step back on the lifecycle of your customer or product, from beginning to end. Design thinking helps in this approach, allowing us to see the interactions and transactions that are most impactful to the relationship.

Get the value proposition right

As Simon Sinek describes, start with "why?" Once we have a more holistic view of the lifecycle of our customer or product, identify the outcomes that are most impactful to improving the experience. Again, design thinking is helpful here.

Value propositions are also shifting from mutually exclusive ROI calculations, into value that is more broadly defined, driven by impact – such as accuracy of estimations, decision-making insight, improvement to customer experience. The individual use cases that will connect to provide this value should be viewed as a portfolio, rather than evaluated and a singular use case level.

Rapid prototyping[34]

The ability to rapidly identify, collect, and model data provides the ability to quickly align between customers, developers, and process owners to improve the customer experience and the effectiveness of the process.

2. Platform

Many industries still have legacy technologies that were originally designed for robustness and reliability, rather than nimble and powerful processing for insights at the edge. The combination of cloud technology stacks and third-party digital software has made this hurdle easier and less expensive to overcome. Technology adoption and migration to the cloud has accelerated significantly over the past five years, with help from the push to work remotely during the Covid-19 crisis. These technology advances have moved to make the platform more robust, flexible, capable, and scalable.

We see that recent improvements to the technology platform have made it less of a concern in scaling programs – only 8% of executives say technology is the limiting factor in moving AI to widespread adoption. And although that may be true, it certainly does not mean technology won't be an important consideration and implementation during the journey.

2x

Organizations are almost twice as likely to report widespread AI adoption when taking a holistic approach vs a piecemeal approach.

VERTICAL AI

To take a holistic approach, zoom out, get the value proposition right, and embrace rapid prototyping.

Accelerate AI Through Platform

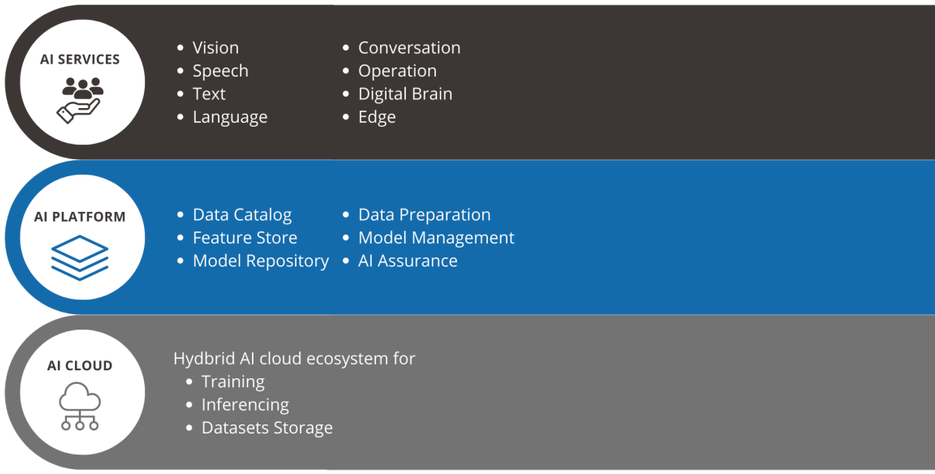

Figure 14: Capabilities across cloud, platform, and services components for AI. Source: Infosys.

However, most organizations are not yet setup in a way to allow the cross-collaboration necessary to achieve vertical AI. Teams sit segmented by function, business process, office (front/middle/back), executive, or other method. To achieve the collaboration necessary, teams should seek to sit cross-functionally across the platform, and from end-to-end across the customer relationship that will be necessary for a vertical AI viewpoint. These teams should own the execution of the zoomed-out strategy noted above, the value proposition and metric identification, the technology assets and data, and the execution components (budgets, schedules, value recognition) to achieve it.

In a similar way to how DevOps enabled more robust, tested, and business-

aligned technology to be delivered more efficiently, MLOps is poised to take AI development and delivery processes to the next level. MLOps provides standardized, automated processes, improves bias detection, and consistency in the approach to handling data and building models. With rapid prototyping and modular readiness to replace models as they develop, we can more efficiently and effectively build the models to support the underlying use cases and more holistic vertical AI value proposition.

3. Enablement

Less than half of organizations – 44% – report that they have well-established policies and practices in place governing data and AI ethics and the responsible use of data. Further, less than a quarter of data executives – 21.6% – believe that the industry has done enough to address data and AI ethics issues and standards.[35]

But at a regulatory level, we're starting to see that change. An AI Index analysis of legislative records on AI in 25 countries shows that the number of bills containing "artificial intelligence" that were passed into law grew from just 1 in 2016 to 18 in 2021.[36] Spain, the United Kingdom, and the United States passed the highest number of AI-related bills in 2021, with each adopting three.[37]

Scale AI Enablement & Innovation

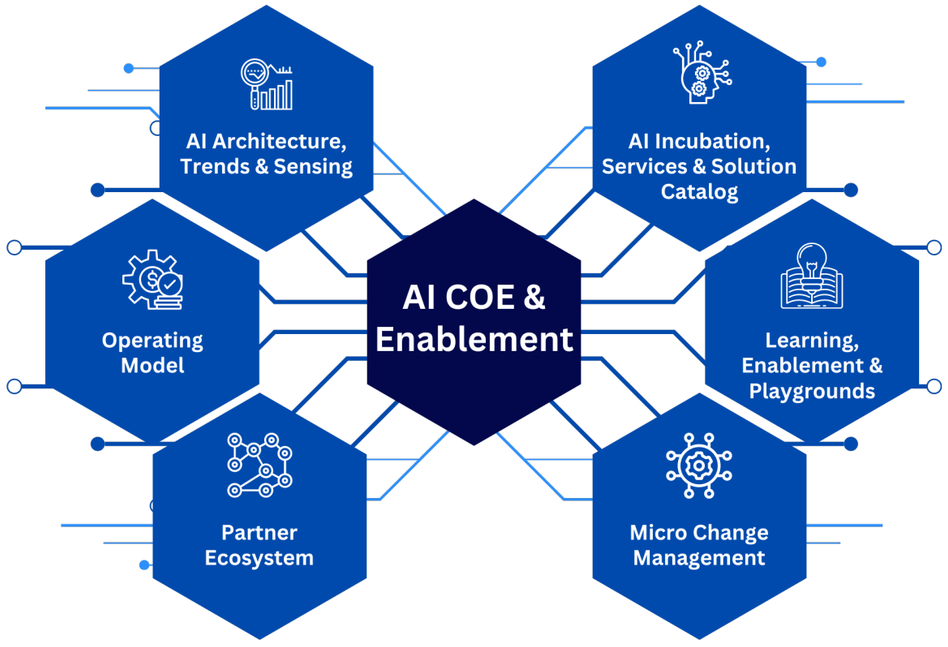

Figure 15: Consideration for scaling AI through an effective an innovative COE. Source: Infosys.

Adoption of responsible AI principles is necessary, particularly as organizations further implement AI within critical business functions. With a vertical AI approach, taking a holistic view across use cases, the goal is to impact critical business functions with real-time data insights.

A comprehensive responsible AI program should take the format of both programmatic and data or model-based requirements. Programmatic elements of responsible AI contain the strategy, operating model, third party, data management, and use case criticality. Model monitoring and explainability are at the core of allowing confidence in consistent and ongoing model output, fairness in its decisioning, and understanding of how the decision was made. Through these governance components, organizations should aim to achieve the following principles around their AI models: purposeful, ethical, human-reviewed, bias detection, explainable, accountable, and reproductive.[38]

AI GUIDING PRINCIPLES

- Purposeful
- Ethical
- Human-reviewed
- Bias detection
- Explainable
- Accountable
- Reproductive

The Next Horizon in AI is Well Within Reach

Vertical AI is an opportunity to "go big," as organizations move beyond limited pilots and siloed use cases. Although organizations have begun to see the potential for AI, in the form of positive returns in most cases, we've yet to see widespread adoption at scale, though a holistic approach and enabling compounding value proposition.

We recommend organizations double down on their efforts to scale AI over 2022 and beyond. The holistic viewpoint across strategy, scaling through the platform (and utilizing the platform as an operating model), while accelerating through enablement will support your organization's success in achieving vertical AI.

And if in doubt, start thinking like a farmer.

About the Author

RYAN MARTIN

Partner, AI & Automation; Boston

Ryan has been leading business transformation programs, utilizing an array of digital technologies, for over 20 years. He focuses on business process improvement through workflow, process automation, artificial intelligence, and machine learning-enabled technologies such as process mining, document processing, and conversational agents. Ryan brings deep industry experience in financial services, and additionally has experience serving clients across technology, consumer goods, travel and airlines, and healthcare sectors. He holds an MBA and BS in Computer and Systems Engineering. Prior to his professional services career, he served as a nuclear submarine officer in the US Navy.

AI PERSONALIZATION

MACHINE LEARNING TO BETTER IDENTIFY PRODUCT & MARKET SYNERGIES

CHRISTOPHER HAYS • SERGE PLATA

Personalization through AI is the next frontier for innovative business analytics. In this article, we look at how businesses can use machine learning to better identify synergies of products and/or markets based on purchase data alone - saving significant time and cost.

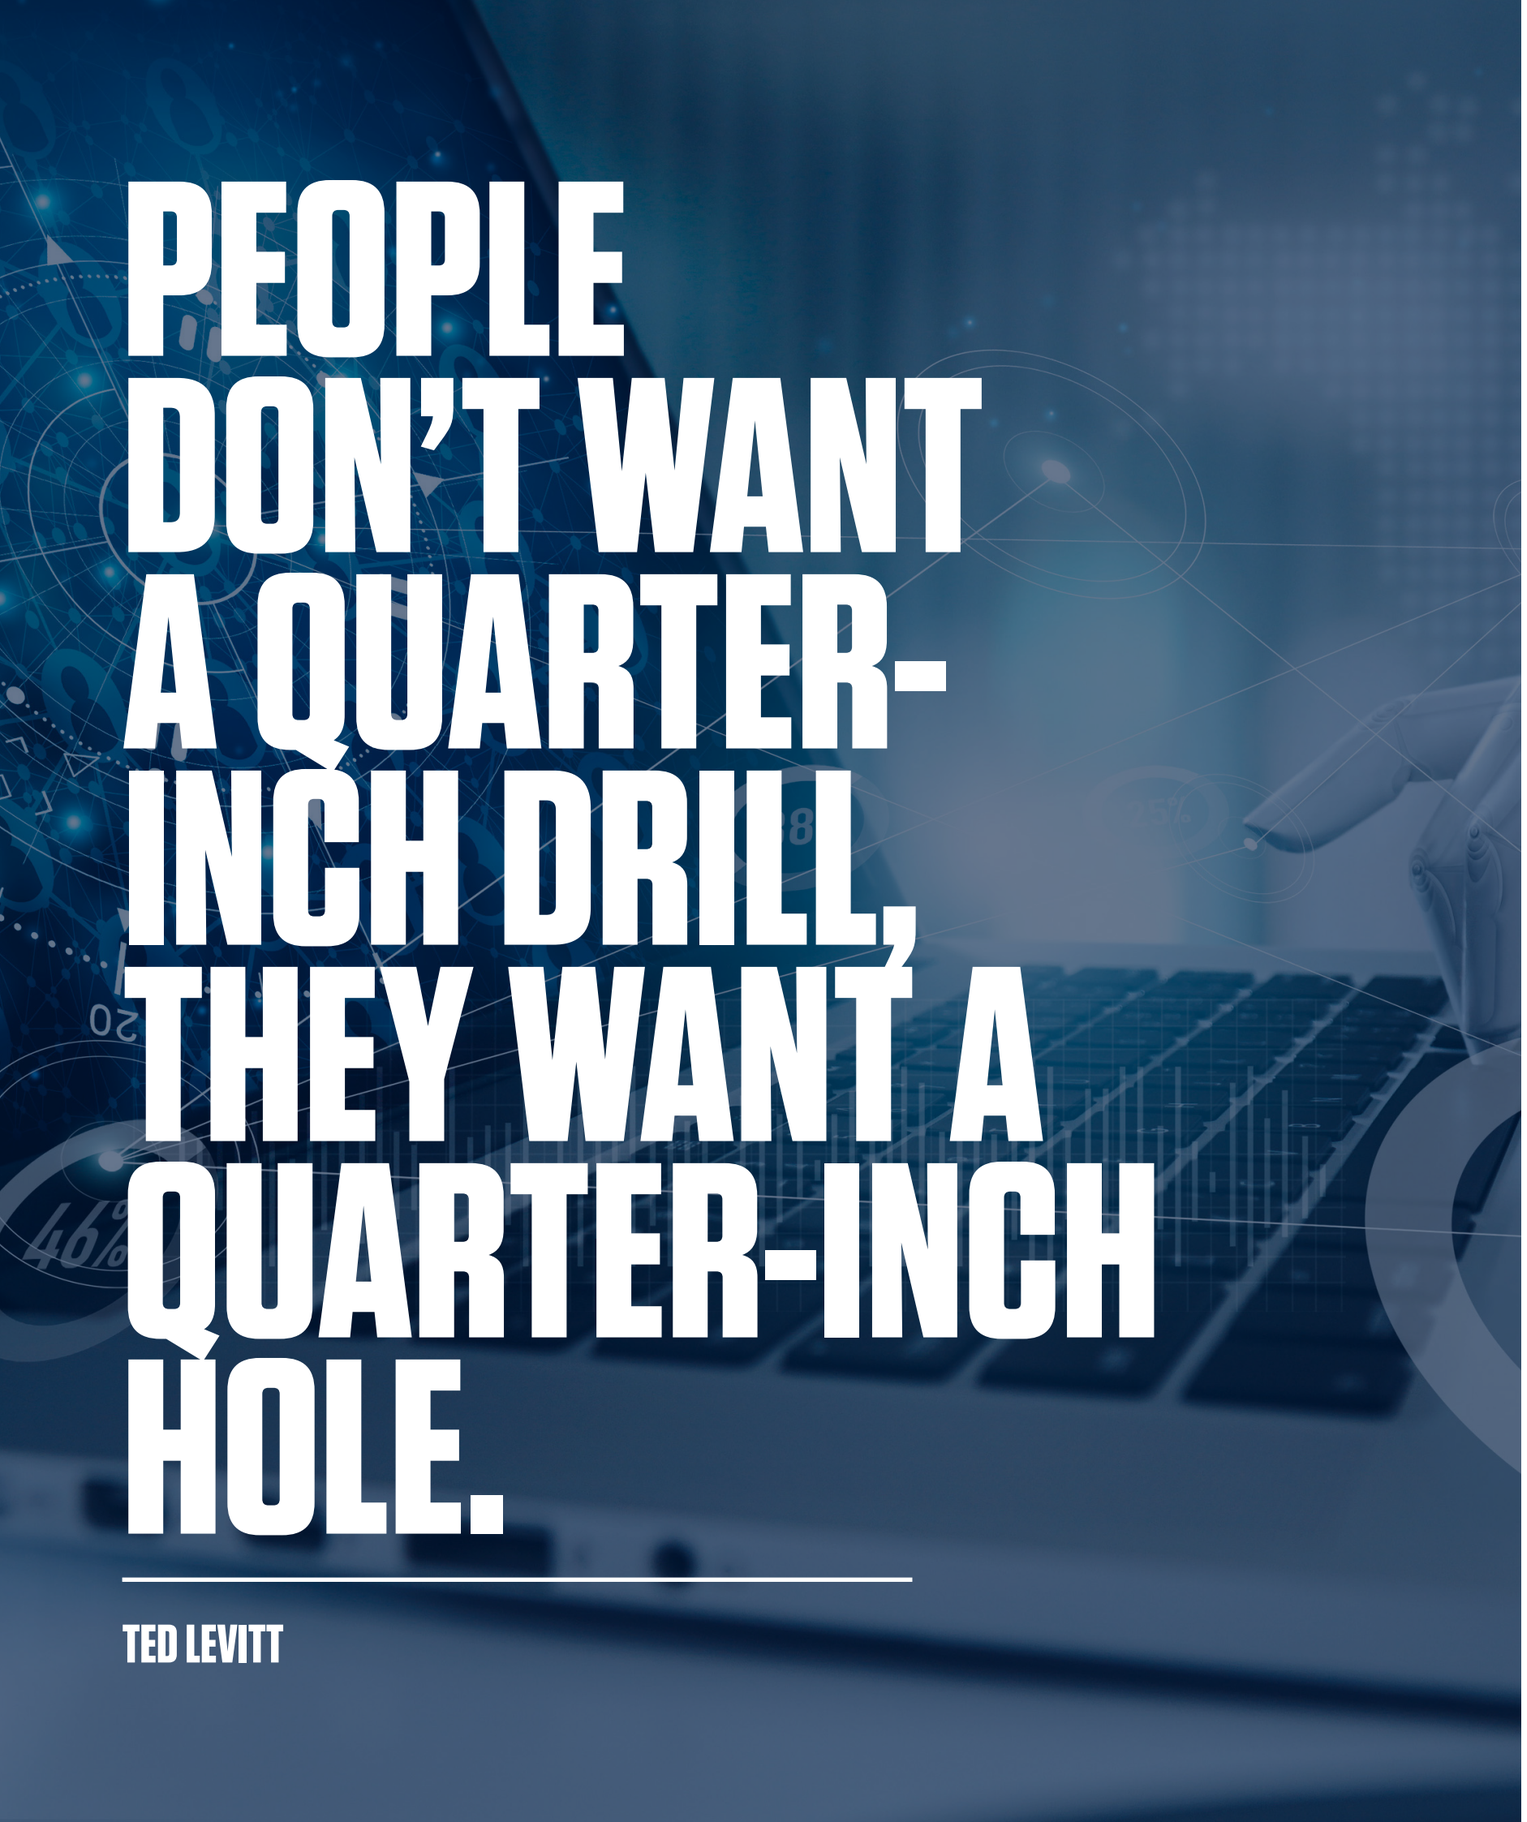

Introduction

Companies interested in reducing competitive threats have focused automation strategies on personalization. Personalization will aid the company in derivative innovation efforts, leaving the threat of disruptive innovation forcing them into obscurity or irrelevance. However, training machine learning (ML) models on purchase data alone can deliver meaningful results that matched customers' missions and revealed new insights only given in cross-category analysis.
This analysis was structured solely on why customers engage (not who the customer is) and unveils the jobs to be done (JTBD) or customer mission.
This analysis output identifies potential unmet customer needs for use by commercial subject matter experts (SME) to identify synergies of markets/products more expeditiously. Combining the identification of the new insights and the confidence in the customer mission fosters an ideal setting for disruptive innovation.[39]

Don't Give Them What They Ask For... Give Them What They Want!

Disruptive theory identifies and predicts the market environment when companies are at risk of being disrupted by a more agile competitor. The theory does not identify the disrupting product or service, define the market potential, give insights on the needs of the customer, or define the competition. That requires a deeper and more holistic understanding of the customers' mission. These customer missions or "jobs" are complex as they include internal and external motivators, as well as the current environment when the job is being done. Most importantly, the actual job is notoriously difficult to elicit, as it is not always transparent to the customer; this is reminiscent of a quote (sometimes attributed to Henry Ford) about customers asking for "faster horses" instead of automobiles. The customer mission can be illuminated when machine learning models are leveraged to review transaction level data across product categories as a product is known by the company it keeps. Using these insights to augment the innovation efforts of SMEs provides speed and focus which can be the difference between disruptive innovation and derivative innovation.

Who is the Customer? (And Does it Matter?)

The difference in this model's approach versus current market models is that it is limited and trained without any information about the customer. Aside from better identification of market gaps, additional benefits readily emerge when focusing on why customers engage versus who the customer is. This method better insulates practitioners from more recent concerns about consumer privacy and inherent biases.[40]

Anecdotally, most of any public unease about artificial intelligence (AI) and machine learning can be attributed to general confusion around AI's capabilities and subtleties.[41] It cannot be ignored that the drive to provide consumers with personalization fueled the collection of consumer-specific data for model inputs.[42] The proliferation of automation – and the general confusion about AI's capabilities – added to the ubiquitous collection of personal information has naturally created consumer demand to know how machine learning affects their lives. These topics speak to deep fears as what interactions are being managed by machine learning and how algorithms could impact academic opportunities, as well as access to creating wealth.[43]

The consumer demand for explainable artificial intelligence (XAI) has created a need for companies to justify their models' decisions (e.g., the same decisions for which they needed to build the machine).[44]

> **The object is to always understand the situation, not the customer.**
>
> Clayton Christensen

Project Inspiration

This article was inspired by the lectures on disruptive innovation by Prof. Clay Christensen presented at Harvard Business School in 2017. According to the theory, by linking machine learning algorithms with disruptive innovation theory to reveal jobs-to-be-done, companies have a tool to innovate. The case was based on a general merchandise retailer in Europe that sells thousands of products, where defining jobs-to-be-done is a concept often difficult to outline. AI and data science techniques were applied to systematically help define the jobs-to-be-done. In addition, given the GDPR and data privacy regulations in Europe, product and transaction data would be used and customer data untouched.

These techniques generated meaningful results that matched customers' missions and revealed new insights based solely on cross-category product and transaction data analysis. Buyers, trading teams, and marketing efforts are normally defined by products [Figure 16], but this model challenges that mindset more than actual jobs-to-be-done.[45] This analysis provides critical guidance, and the output should be used to enhance experts in the field and multidisciplinary teams to simplify and synthesize the analysis.

The Most Frequently Purchased Products on Amazon

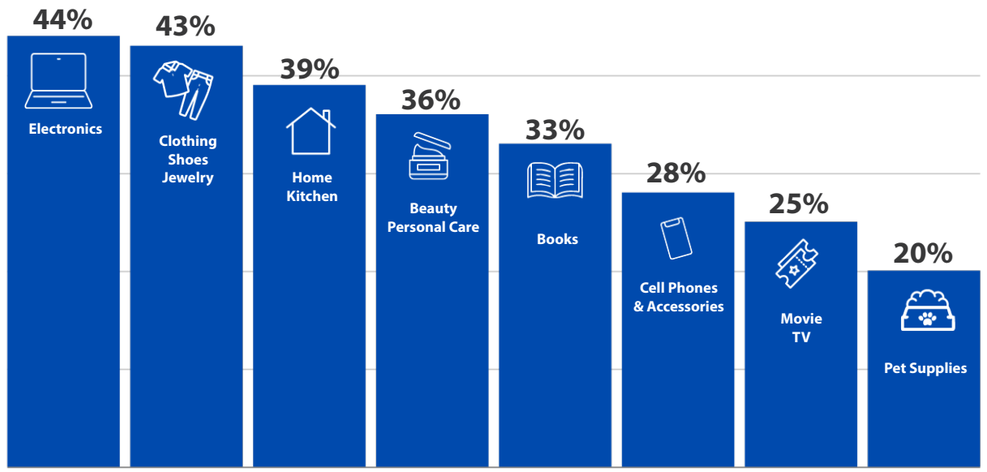

Figure 16: Most firms define markets by product category. Source: https://financesonline.com/amazon-statistics/ (Accessed on 06/28/22) (Feedvisor 2019).

Innovation: Disruptive or Derivative?

One of the main factors in the disruptive innovation theory is the concept of jobs-to-be-done. Narrowing down the job or jobs a company is helping its customers to do sounds easy, but it is difficult, especially when the case at hand is a general merchandise (GM) retailer or a supermarket due to the variety of products available. The fundamental problem is that most of the masses of customer data that companies create is structured to show correlations. Purchase and transaction data is often structured in a way that results in a self-fulfilling prophecy where relationships and conclusions seem only to confirm what is already understood.

In disruptive innovation, one must think about data differently, and outlining jobs-to-be-done is a great way to start.

Traditionally in retail, markets are defined by products; for example, establish the "mobile phone market" or the "garden furniture market," but when trying to apply the disruptive innovation theory, markets are better defined by the job-to-be-done. So, it is not about the specific products but about customer needs and behavior.

"Gazillions" of Product Combinations

In the case of general merchandise retailers, tasks like outlining a promotional activity or developing a sales strategy normally require multiple experts from different areas, not only in data but also subject matter experts who analyze all possible outcomes. It is harder and harder to do these analyses as companies serve more customers with a wider variety of products and subproducts, while their ecommerce branches reach many markets and demographics at the same time.

A common task for data scientists, and one of the main challenges encountered, is dealing with copious amounts of data and simplifying the findings for the business' SMEs to make better decisions. Sometimes analysts and data scientists must deal with thousands of products and thoroughly analyze all the possible avenues of commercial and logistics decisions.

So, the approach was not to focus on a product but a set of products that define a certain mission or a job-to-be-done. Sometimes these "missions" are well defined; for example, "buying a toy" might be related to a birthday present or Christmas. On the other hand, if that same toy is bought together with a stroller, then the job might not be to "buy a present for a birthday" but to "raise a child." On many occasions, the missions are so intricate that they are very difficult

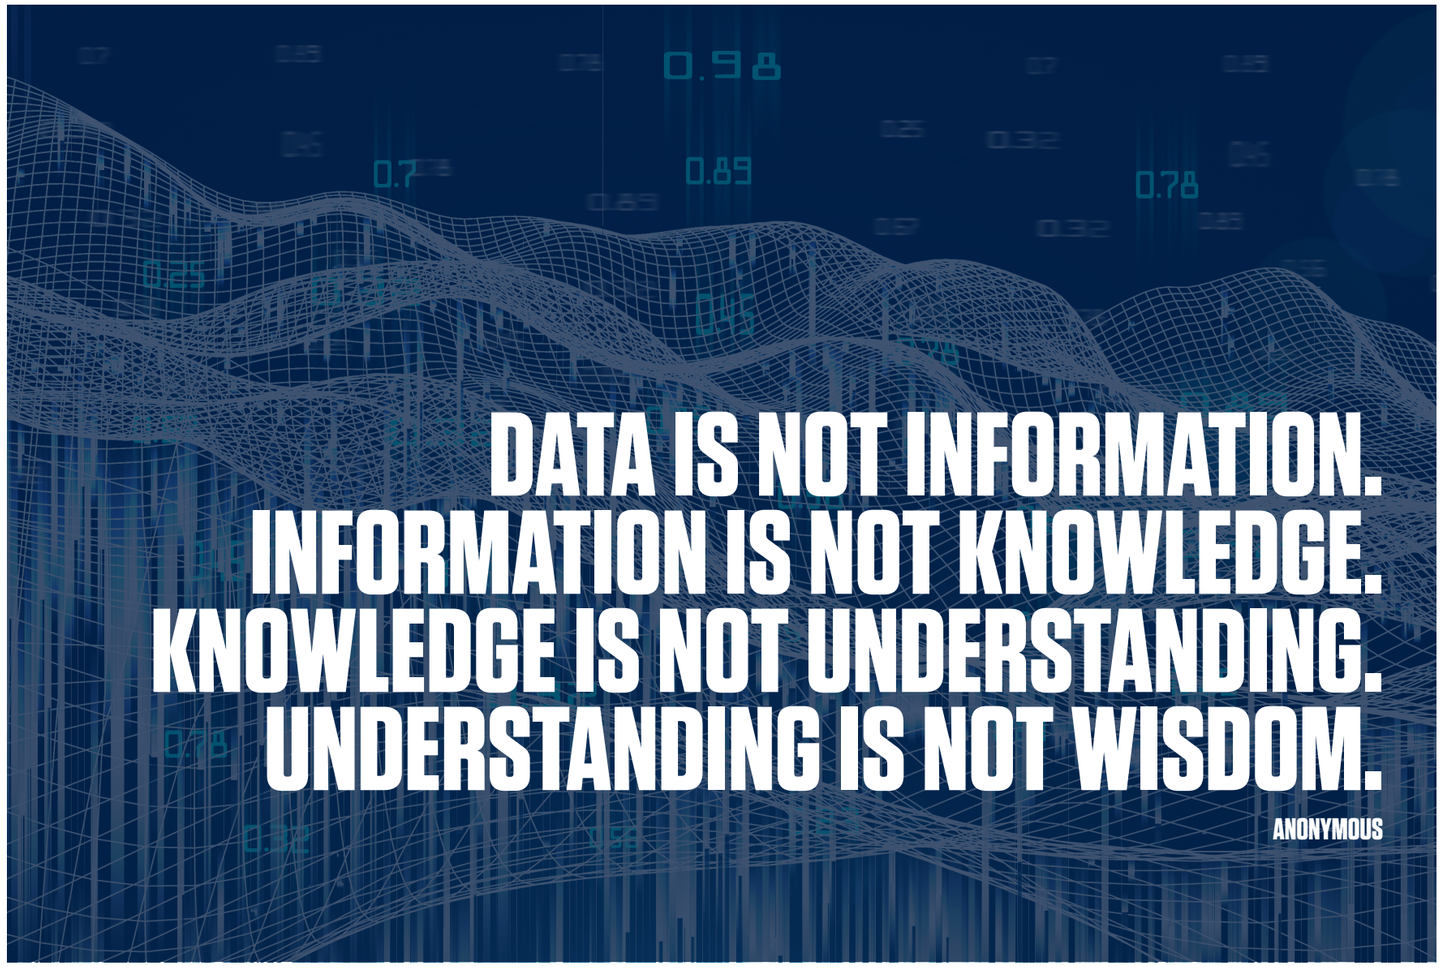

> **In the factory we make cosmetics. In the store, we sell hope.**
>
> **Charles Revlon**

to recognize, not to mention experts in differing specific markets creating convoluted involvement when identifying overall customer need.

Analyzing these bundles of products can be time-consuming, and if one wants to be thorough about it, analyzing all the possible combinations can be exhausting. For instance, if a company has 1,000 products in stock for sale, all the possible combinations of pairwise bundles is $2^{1000} = 1$ followed by 301 zeroes or a novemnonagintillion combinations. The incomprehensibility of large numbers illustrates the inherent need for a system to raise these combinations to human experts for review.

(Machine)/ Learning the *Why* from the *What*

From the data science perspective, a systematic way was created to identify the JTBD using transaction data only. The problem mainly relied on the fact

that there are thousands of products in stock available for customers and almost never a way to pinpoint behaviors. Reasons for this dilemma range from customers changing their behavior quickly to the number of product combinations quickly exceeding current computational limitations.

This is especially difficult to do when retailers sell across different demographics and geographic zones in the entire country. Relating certain behaviors to a job-to-be-done could be achieved by overlaying data from different sources, for example, weather, price, time of year, and quantity of product bought. This complicates the already complex task of bundling products.

Simplifying the analysis did little to mitigate the proliferating variables. Jobs-to-be-done, linked to products, might have had many underlying reasons: price, availability, taste, or consumers' lifestyle. So, in general, there is a subject matter expert at many retailers with a lot of experience in assessing, mainly by intuition, what sales could have been linked to a basket or–by previous experience–pointing to a certain "mission." Also, companies that seek to address the fundamental problem the customer is trying to solve typically expand the applicable market.

Bottomline, it is about recognizing *why* people are buying **what** they are buying.

Methodology

Machine learning has been around for a long time since IBM's Arthur Samuel's checker game in 1952. Deep learning is a subset of ML, which is a subset of artificial intelligence, all of which are far from completed topics with steady advances in all areas.

The homogeneity of products is useful for a retailer when grouping similar items. Unfortunately, this can lead to a frustrating experience for product recommendations. A single sale of a kettle does not signal the inaugural purchase of a new collection of kettles.[46] Selecting items for customer missions or jobs-to-be-done requires a heterogeneous mixture of linked products.

A direct, simple solution would be to match every item with every other item. While this guarantees the discovery of cross-category relations, the search space is massive, and it grows exponentially. With products in the hundreds of thousands, this approach is impractical.

Instead, embedding algorithms and vectorizing products helped us to standardize the dataset; following[47] a semi-supervised algorithm, surfaced latent patterns in sales data. Using the representations this self-learning algorithm learned a systematic similarity measure, thereby removing the need for crafting features by hand. Suitable products for the jobs-to-be-done campaign were now identifiable. This enabled a methodology for companies to focus on disruptive innovation. Given enough data, the quality of the data ceased to be an issue. Python, the lingua franca of industrial data science,[48] had no issues processing data at this scale.

> **There is no single right answer or path forward, but there is one right way to frame the question.**
>
> **Clayton Christensen**

Known links between products served to verify the associations revealed by the algorithm; for example, a paddling pool listed closely to a table fan. The hypothesis was that the algorithm would reveal previously unknown links between products. Promotions targeting customer missions could use these links and would have existing sales as their foundation.

The idea of customer missions or jobs-to-be-done is to link related products with statistical significance. The goal of the algorithm is only to provide candidate links for teams to consider. It is not intended to replace the human team but to empower them by putting products in bundles that address customers' needs.

Algorithm

In traditional data science projects, transactional data by itself does not give clear insight into buying patterns. The data needs annotations, enrichment, and cleaning. This data can originate from ecommerce, surveys, focus groups, or even internal metrics. However, these sources are inconsistent, needing care and attention whilst pre-processing. To respect customers' privacy, this type of enrichment was not performed, circumventing unnecessary pre-processing.

Finding a balance between data that was not customer-specific but exposed customer patterns proved difficult. For instance, overlaying in-stock/on-shelf data could provide insights into patterns. However, this would not definitively prove the existence (or absence) of a link between products.

A Product is Known by the Company It Keeps

The algorithm used was a neural network and created an index or measure for products based on the items sold at a similar time or location. This resulted in listing similar items that were either complementary (batteries with battery-operated toys) or substitutes (different brands of the same product).

The analysis showed that it can be deceiving to use closeness of these representations as a measure of similarity. This could be due to the price, availability, immediacy, or other factors independent of the customer's mission. Spotting these patterns or customer missions requires domain knowledge with a plan to automate as much as possible.

Combining these ideas resulted in a pipeline where the algorithm's representations of products were processed using other techniques like k-Nearest Neighbors. This combination has seen success in natural language processing (NLP), a subfield of AI focused on analyzing natural language data. Adapting an open algorithm like the one used in this case for retail purposes has precedence in Biswas, 2017 and McCormick, 2018. However, the use detailed here differs in that it focuses on customer behavior instead of product attributes. This difference is key in leveraging the success of the algorithm in tasks like sentiment analysis and focusing it on customer missions or jobs-to-be-done.

Hurdles

The main hurdles in the analysis were: 1) data and 2) product/market structure.

Data – importance of the right data volume

As with most companies, too much or too little data hinders proper analysis. When there isn't enough data, it is hard to generalize analyses. If there is an excess of data, noise can mask the signal (the goal) from detection. Processing such large volumes adds computational costs, pushing budgetary limits. Apart from volume, data quality needs careful consideration to avoid misleading analyses. Finally, GDPR[49] and privacy regulations are getting companies to rethink the usage of data. Finding a balance in the volume, type of data, and data privacy regulations along with its quality was crucial.

Product/market structure

Hierarchical organizational structures result in teams owning sections of the product or market. Cross-pollination of ideas and talent is imperative as, over time, this blinkered approach leads to talent silos. These stand-alone repositories of expertise do not allow for organic cross-team collaborations. For example, a cross-team collaboration approach would have found the aforementioned connection between table fans and paddling pools. Both are popular choices for staying cool during a heat wave. And yet since these products usually fall under different markets, the customer mission is further obfuscated, additionally complicating addressing those needs.

Case Studies 1& 2: Connecting the Dots for Better Product Recommendations

Case 1: Algorithm associated a budget TV with a budget laptop, weighing scale, and phone.

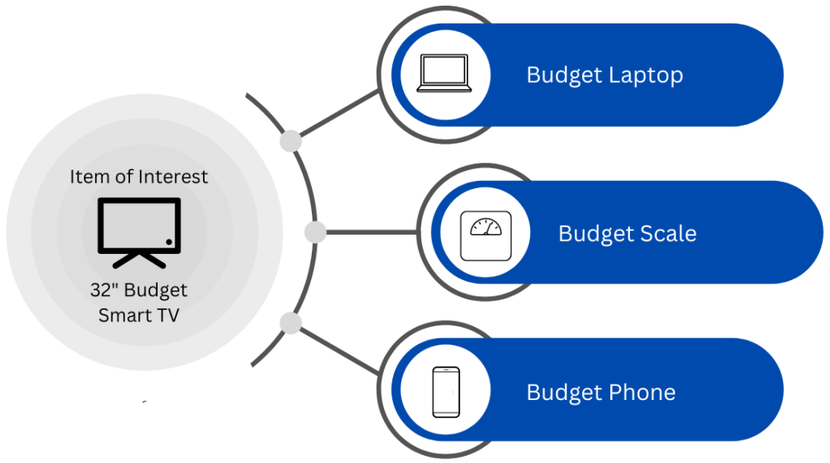

Figure 17: Algorithmic product recommendations based on budget TV interest.

Case 2: Algorithm associated a luxury TV with a luxury TV, a microwave, and an iPad.

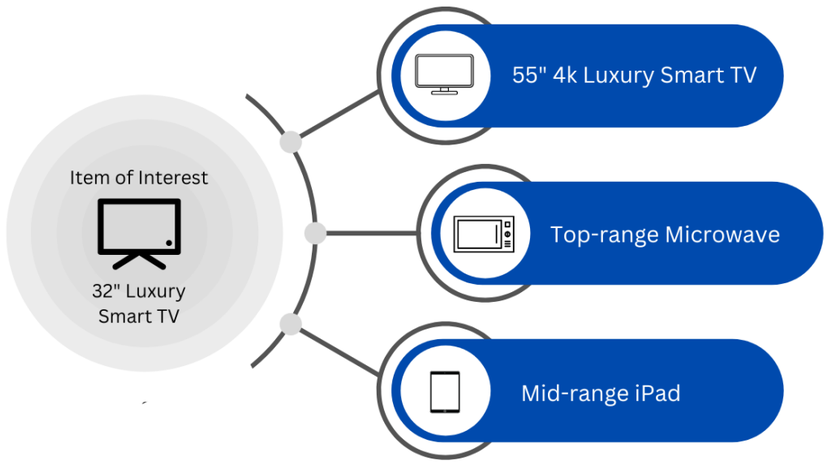

Figure 18: Algorithmic product recommendations based on luxury TV interest.

Effective Promotions Should Address the Overall Needs, Not the Market-Specific Ones

Results

An interesting example of the algorithm's findings: the input of two TVs, one a budget and the other luxury. Previous derivative algorithms trained on only a single product category would have output other similar TVs in both cases. This model's algorithm was able to "join the dots" cross-product and correlate seemingly unrelated products to find disruptive possibilities. Analyzing these results was very reassuring, as clear missions could be associated with certain products: in the first case, the mission was to furnish a home with "basics," for example a landlord trying to rent a property. In the second case, it was also furnishing, but with the corresponding associated products. This certainly outlines a lifestyle more than only associated similar products or randomly clustering them.

This approach is transferable to other products, for example, the below Tweet posted by Medium about the correlation between produce and insurance. More importantly, the reach of this approach can shed light into the customer's lifetime value as it follows trends in time.

Medium
@Medium · Follow
Purchasers of fresh fennel are linked to low insurance claim rates, and other insights data brokers are finding off of information you may not have even known was being collected read.medium.com/HodN1ZN
2:43 AM · Sep 18, 2018

Figure 19: Tweet from Medium, Sept 18, 2018

In Case 3, the mission associated with the game input was "to raise a child" rather than "just getting a game." The model output a half-size guitar and headphones go with it and reinforce the conclusion, as well as similar games for a certain age.

Case 4 was unique: the first product that the algorithm associated with the input of a toddler toy was a specific tumble dryer accessory that speeds up the drying process. In the beginning, it appeared that there was an error in the program. After consultation with commercial SMEs and examining anecdotal surveys, the results were sensical and explainable as a household with toddlers often must

Case Studies 3 & 4: Recommended Products Taking Into Consideration Age of Child

Case 3: Algorithm associated an older child's video game with a junior guitar, headphones, and age-appropriate ages.

Figure 20: Algorithmic product recommendations based on video game interest.

Case 4: Algorithm associated a toddler toy with a tumble dry accessory

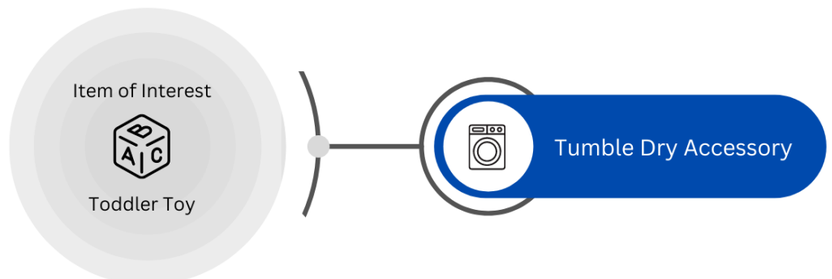

Figure 21: Algorithmic product recommendations based on toddler toy interest.

> A strategy is nothing but good intentions unless it's effectively implemented.
>
> Clayton Christensen

wash more than a household without toddlers. Subsequent tests confirmed the two products were indeed statistically close in sales as the algorithm identified.

The association of these two products was great news as the methodology and algorithm identified the cross-product pattern that would be very difficult for SMEs to identify. SMEs are often confined in their own areas of expertise (or markets) and cross-category links are very difficult to catch.

Conclusion

This article demonstrated the links between AI and a very important factor in disruptive innovation theory: jobs-to-be-done. The work is still in the research and development phase. It does demonstrate that meaningful results matching customers' missions or jobs-to-be-done can be produced using machine learning algorithms trained on purchase data alone. Given this method's inherent privacy, it can help to support innovation theories in a systematic way and lead to better support for disruptors. For example, use of these methods by commercial subject matter experts can help better identify synergies of products and/or markets.

About the Authors

SERGE PLATA

Senior Principal, AI & Automation; London

Serge is a certified six-sigma black belt, and a chartered scientist in the UK. He was awarded a fellowship from the UK Institute of Mathematics and has more than 15 years' experience building, leading and consulting for data programs from FTSE-100 companies to SMEs and start-ups. He is currently the Chief Decision Scientist for Infosys Consulting EMEA.

CHRISTOPHER HAYS

Principal, AI & Automation; Alabama

Christopher has been leading digital transformations for the past twenty years, with the most recent fifteen years in transportation and logistics. He has practical experience solving complex issues by combining his experience in design thinking, change management, and process design with technologies in AI & analytics, IoT & data solutions, as well as prediction & optimization methods. Solving these complex issues for his clients drive operational efficiencies, reduce costs, identify strategic opportunities, and ensure overall value realization.

BUSINESS FORECASTING

A NEW APPROACH TO FORECAST ACCURACY

Improving Forecast Accuracy to Significantly Increase Business Value

SERGE PLATA • MAGDA RACZKIEWICZ • HAIDONG (HENRY) SUN

Forecasting is crucial in consumer goods, but there is often a disconnect between metrics and actual business needs. In this article, we look at a better way for data scientists to evaluate forecasts for improved business outcomes.

Introduction

There is a common issue in business problems, particularly with data-led projects: do internal metrics of the model align with the metrics of the business? For example, if a retailer increases forecast accuracy by 1% would they move less boxes in the warehouse, increasing efficiency and therefore experience elevated profits? Or is the metric disconnected from the actual business operating model?

This point of view analyzes accuracy metrics in predictive models and their relationship with other metrics that are closer to the business. There are several studies (like Hyndman, 2006) that analyze forecast metrics, and recommend the use of MAPE, MdAPE, sMAPE, sMdMAPE, MdRAE, GMRAE and MASE. However, these metrics could mislead conclusions and business outcomes.

Indeed, many companies use MAPE (Mean Absolute Percentage Error) to evaluate forecast accuracy as it is easy to explain and communicate. However, this metric has some limitations and - as we will show - sometimes conceals real issues in supply chain and other commercial activities. MAPE or RMSE (Square Root of the Mean Square Error) do not reflect concrete actions to increase productivity, efficiency, or operational improvements in organizations, but rather create confusion - especially when trying to find an action out of the analysis.

We will explore more the differences between metrics and how they relate to real life business situations. We believe that the performance of forecast accuracy metrics should be linked to data, models, and most importantly to business actions and business decisions behind the process.

A Compelling Example in Retail Forecasting

We analyzed a case with a retailer in Europe and a baby store in the U.S. Both retailers had several types of products, from slow to fast movers, seasonal, stable, and volatile products (see Important Terms box, next page).

The retailers have a strong forecasting department, since this is crucial for the operation; sometimes the delivery lead time is long (up to three months), so reordering is out of the question. They must get it right the first time.

Forecasting was measured with MAPE and efforts to improve the forecast were made based on this metric. One of the most important questions was "if we increase the model accuracy by 1%, how would that benefit the business?" However, this was very difficult to answer as MAPE had some limitations and concealed issues like volatility or bias.

To answer this question, we proposed establishing a close link between the internal metrics of the model and other external metrics from the business, like sales, or actions towards a better inventory planning.

We found that GSMRE (Geometric Squared Mean Root Error)[50] is a better metric than the standard MAPE[51] or RMSE for certain operational situations, for example when dealing with slow movers, as it correlates well with metrics measured in units like MAE (Mean Absolute Error) more than a percentage.

In addition, one of the most important outcomes of this article, on top of the actual analysis, is the implementation and calculation of metrics like GSMRE, which might easily overflow computer systems. For this we implemented an iterative function, as recursive functions had an upper bound.

Metrics Review

Practitioners are normally biased when choosing forecasting or predicting models, depending on the type of activity or business. For example, a slow mover, like a microwave oven or a set of frying pans would tend to be over-forecasted, which might meaningfully impact stock levels for a warehouse and decrease the manager's KPIs.[52]

Thus, it is important to consider bias (i.e., if the model is over forecasting or under forecasting) while planning the entire evaluation process; in this way, the technical teams may know how to improve the model.

Naturally, the use of asymmetric measures used as loss functions by the data science teams like MAPE or MAE, MSE, or RMSE can lead to biased forecasts (for example over-forecasting or under-forecasting). Several studies have treated the prediction bias as a crucial feature of the forecast analysis.
Often, the RMSE is preferred to the MSE as it is on the same units (or

IMPORTANT TERMS

Fast movers - items that sell fewer than twenty units in a calendar year.

FMCG - Fast-moving consumer goods are products that are sold quickly and at a relatively low cost. For example, packaged goods, beverages, toiletries, etc. promotional items subject to regular promotions during the year. These items show spikes above the 2-σ, but only in specific periods and due to promotional activity.

Seasonal items - sell only once a year, for specific events e.g., holidays, back to school, etc.

Slow movers - Items that sell fewer than twenty units in all stores in a calendar year.

Stable items have low standard deviation. Normally these items will have 90% of the data between $\mu - \sigma$ and $\mu + \sigma$.

Volatile items - Following (Hull, 2005, 466), items with a high standard deviation. Normally these items will fluctuate more than once above 2σ and at least once above 3σ.

scale) as the original data. These two metrics have been popular, largely because of their theoretical relevance in statistical modeling. However, they are more sensitive to outliers than MAE (Mean Absolute Error) which has led some authors to recommend against their use in forecast accuracy evaluation.

Not getting a proper measure can lead to serious problems like over-stocking or running out of products in the middle of the season. For example, if we systematically over-forecast, then the system will receive the signal to buy more, causing overflowing of product.

So, what metric should we use? This would depend on the data and the problem we are trying to solve.

Data & Problem Setting

One of the main problems, often taken for granted, is the choice of accuracy metric. As stated by Kourentzes, "Although most of the literature agrees with the fact that experts to a certain extent improve forecasting accuracy, the results are inconclusive."

Conclusions and recommendations to business units depend on the error metric chosen. For instance, there have been studies where the adjustments improved the forecasting accuracy when Median Absolute Percentage Error (MdAPE) was used, while the opposite conclusion was reached with other metrics. This is of the utmost importance, as the one metric was suggesting the business was thriving and the other totally the opposite!

We decided to analyze some of the metrics and link the results to the type of data and models to assess them and propose a robust and rigorous approach that improves operation efficiency.

A more objective way to assess this is by stating that the relevance of the metric must be linked to the absolute error, which is calculated in the same units as the real objects of the prediction.

For example, if we miss the forecast by 10 boxes, we need to physically adjust 10 boxes in the warehouse, deliveries, store allocation and transport process. Someone must physically move them and deal with them if they do not sell them. That is why the model metrics should correlate.

Models

We chose to apply a well-known and widely used model: ARIMA (Autoregressive Integrated Moving Average). This model is among the most effective approaches for analyzing time series data.

There are other classic and well-performing predictive models used to make forecasts based on time series data, like exponential smoothing. This technique was proposed in the late 1950s and early 1960s or PROPHET, which is based on Fourier transform.[53] The model in this case is not the objective of the paper but the evaluation of the models. The metrics above mentioned are widely used in all cases.

Metrics & Analysis

Using arithmetic mean or geometric mean has pros and cons. In our analysis, metrics that used geometric mean performed equal or better versus others that use arithmetic mean. We tested the metrics with different types of products. The data for the different product types are shown graphically below:
One of the main comparisons in our analysis was a volatile product and a slow mover as shown below. The charts show actuals and forecast for 28 weeks. We

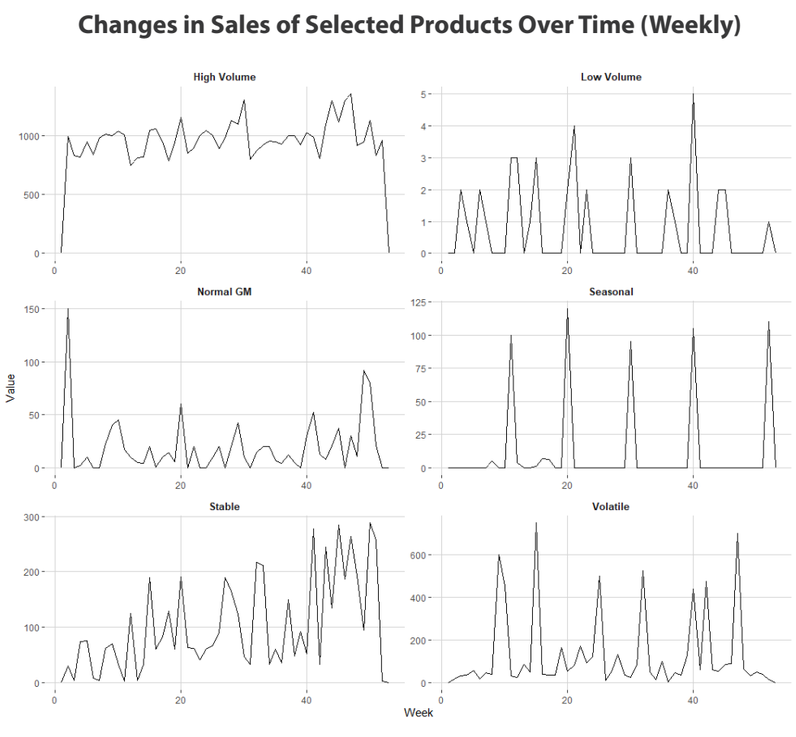

Figure 22: Changes in sales of selected products over weeks.

calculated the MAPE and the GMRSE for comparison and observed that MAPE had a very similar score in both cases, where clearly the forecast was better with the slow mover (almost overlapping with actuals).

	MAPE	GSMRE
Sample Product - Stable Medium Variance	0.62	3.67
Sample Product - Slow Mover High Variance	0.65	0.59

Figure 23: Comparison of MAPE and GSMRE for stable versus slow mover products.

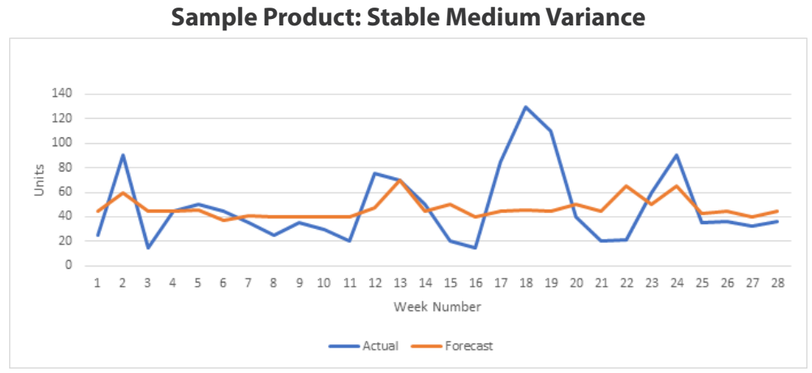

Figure 24: Comparison of forecast vs. actual sales of a stable product with medium variance over 28 weeks.

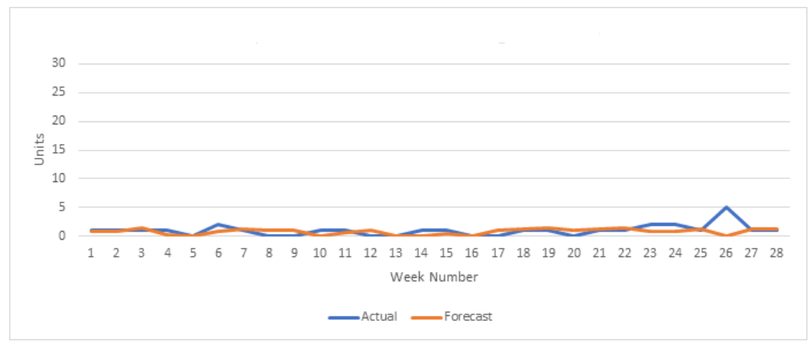

Figure 25: Comparison of forecast vs. actual sales of a slow mover with high variance product over 28 weeks.

In this case, we can conclude that MAPE hides volatility and does not reflect the real differences between products whereas GMRSE gives an index of fit between forecast and actuals across weeks. Also–according to specialized literature–it deals better with outliers, something that we could observe in our analysis too.

The table below shows the results of other metrics.

	MAD	MAPE	MSE	RMSE	GSMRE	Bias Coeff
Stable	5.12	62.18	44.16	6.65	3.67	-0.21
Slow Mover	0.83	64.70	2.28	1.51	0.59	-0.11

Figure 26: Additional metrics from our analysis of the stable and slow mover products.

The correlation between MAD (which is the average of the actual error in units) showed positive with GMRSE and negative with MAPE.

In this sense, GMRSE is a better metric than MAPE as it has direct impact on the actual units. The benefits can be better illustrated by following up on the question above: How many extra staff should we need to schedule to move 200 extra boxes due next month? If the forecast accuracy is low (there is a huge error), then the business might incur reprocessing and inefficiencies that translate into financial results. A common question for the operations managers is how much money we will save if we invest in improving the forecast accuracy by 1%.

If the forecast is measured by MAPE as the example above, then there is no correlation between an improvement in forecast and the actual operation.

If we use GMRSE as in the case above, the metric translates immediately into actions that can be taken by the business due to the positive correlation with the actual units to be moved.

If forecast accuracy is low then the business might incur reprocessing and inefficiencies that translate into financial results.

Results

General results are shown below.

Measure	MAD	rMAD	MAPE	MSE	GSMRE
Units	Units	%	%	Units squared	Index
Stability	Stable	Not defined when actuals are zero	Not defined when actuals are zero	Biased when the difference is large	Stable
Sensitive to	Does not identify bias	Small denominator	Small denominator	Outliers	Large Variations
Cost/Scale trend	Linear	Linear	No link	Quadratic	Logarithmic
Outlier Detection	Medium	Medium	Low	Very Good	Low
Best Estimator	Median	Median	Mean	Mean	Median
Benefits	Measures actual units	Better correlated to MAD than MAPE	Good at aggregate level	Very good handling small differences	Very good at granular level/ excellent correlation

Figure 27: Overview of forecasting metrics conclusions.

CASE STUDY

GSMRE In Action

StableCo is a global direct-to-consumer company with multiple product lines including cosmetics, personal care, clothing, and accessories. It produces a new product catalog every two weeks, and this catalog production is called a campaign; there are 26 campaigns a year.

Thousands of independent sales representatives will place orders for themselves, and their network of customers, based on the campaign catalogs. The company handles order fulfillment and other operations-related functions.

Demand forecast is a key function as forecasting accuracy directly impacts the company's profitability. This is partially driven by the need to have a healthy product churn between campaigns to generate excitement and sometimes long in-bound logistic cycles.

The executive leadership team are evaluating a new product line: baby products. It seems to be a good fit as most of the representatives and customers are females, and many have infants and small children at home. Also, some previous campaigns had baby products with reasonable sales, although the product range was limited. Conducting a traditional new product introduction market research with typical field work is both costly and time-consuming.

Leveraging our expertise in this area, we devised a practical and cost-effective approach coupling forecasting model analysis with scaled-down field work. For baby wear and baby socks, which were consistently offered, we could forecast based on past sales. For other products, such as bath toys, mobiles, and other baby accessories, limited market research was used. As this is not critical to our forecast model discussion, we'll not go into any detail, although suffice to say that it had cost containment measures with scope controls in markets, products, and samples. Figure 28 shows the historical unit sales.

	Babywear			Babywear by Size		Socks		
	Male	Female	Total	Up to 1T	18M-2T	Male	Female	Total
Feb	4,721	7,671	**12,393**	5,909	1,895	15,452	15,681	**31,133**
March	6,442	8,918	**15,360**	9,021	3,785	23,621	16,687	**40,307**
April	8,036	9,956	**17,992**	7,484	8,168	16,819	14,348	**31,167**
May	7,724	9,901	**17,625**	9,018	5,346	13,917	8,840	**22,757**
June	9,084	13,341	**22,425**	12,550	4,784	9,887	11,298	**21,185**
July	7,540	7,997	**15,538**	8,503	3,607	9,945	8,011	**17,956**
Aug	4,723	9,043	**13,765**	8,562	3,006	12,044	10,174	**22,219**
Sept	7,632	11,797	**19,429**	10,314	6,079	15,660	16,512	**32,172**
Oct	6,683	10,626	**17,308**	10,141	4,594	25,782	14,105	**39,887**
Nov	8,399	7,006	**15,405**	7,930	5,749	20,579	13,114	**33,693**
Dec	10,858	18,864	**29,722**	11,019	14,138	27,972	21,974	**49,946**
Jan	5,704	6,939	**12,643**	8,115	2,947	11,969	15,152	**27,121**
Feb	5,971	7,617	**13,588**	8,215	3,340	13,059	14,444	**27,503**
Mar	13,331	8,719	**22,050**	10,243	6,561	25,447	19,429	**44,876**
April	5,991	10,203	**16,194**	10,732	3,336	12,139	17,420	**29,560**
May	9,938	12,303	**22,240**	15,898	3,989	21,842	9,318	**31,160**
Jun	10,089	11,266	**21,355**	10,531	7,470	22,631	12,820	**35,451**
Jul	9,384	10,096	**19,479**	11,303	5,032	12,465	8,489	**20,954**

Figure 28: Historical sales of babywear at StableCo.

The ARIMA forecast on baby wear and baby socks showed the following charts:

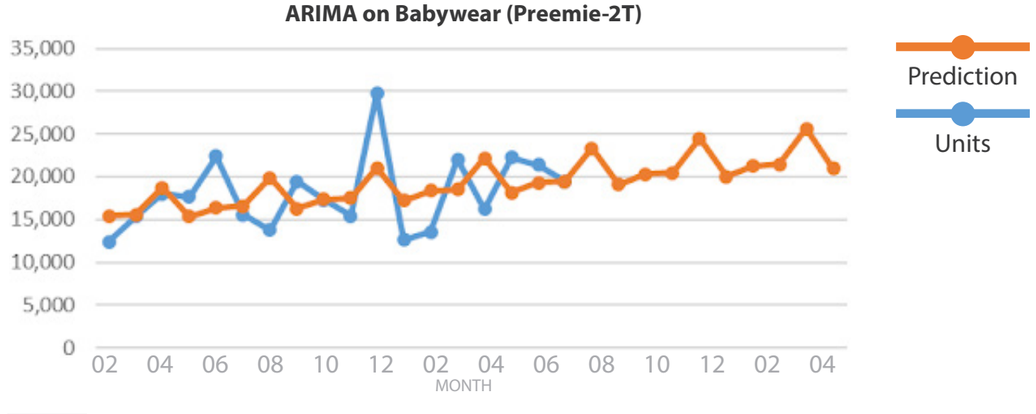

Figure 29: ARIMA forecast for babywear.

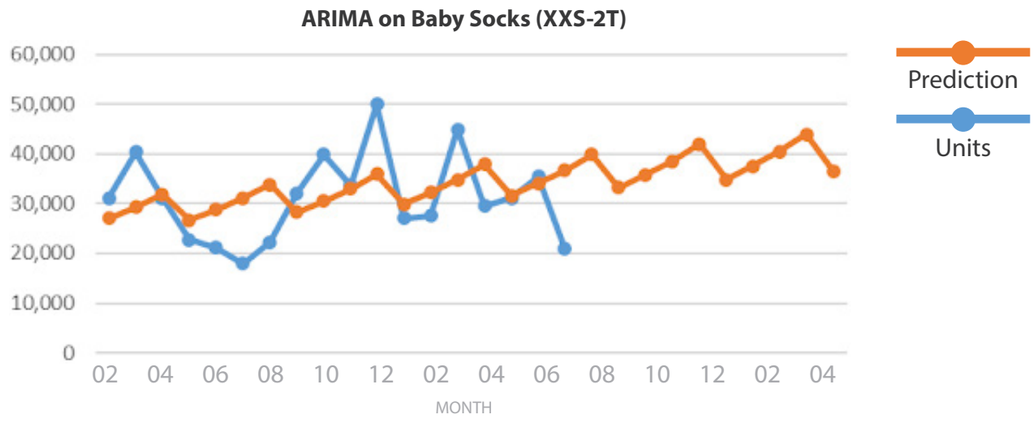

Figure 30: ARIMA forecast for baby socks.

One can see the model captured seasonality and trend, however the forecast versus the actual is not as accurate when dealing with the numbers point by point. We would like to deal with the numbers month by month because that is what the operation requires.

The following table shows the comparison of accuracy metrics for different product segments:

	MAD	MAPE	MSE	RMSE	GSMRE	Bias Coeff
Babywear Male	1,739	23	3,664,893	1,914	1,518	-0.11
Babywear Female	2,007	20	6,181,446	2,486	1,353	0.22
Babywear Total	**3,243**	**18**	**16,233,115**	**4,029**	**1,562**	**-0.22**
0-1T	1,501	16	3,774,642	1,943	1,029	0.00
18M-2T	1,878	41	6,214,118	2,493	1,403	0.00
Socks Male	4,952	33	34,416,316	5,867	3,422	0.11
Socks Female	3,016	26	13,941,092	3,734	1,657	-0.11
Socks Total	**6,875**	**25**	**70,756,913**	**8,412**	**4,497**	**-0.11**

Figure 31: Comparison accuracy metrics for different product segments.

As one can see from the table, MAD and MAPE do not correlate: for example, increasing the error in units between Babywear male and Babywear female but decreasing MAPE. Errors in units are costly and the forecasting model is of little use if it is not improved. The problem is that the model's error (and therefore accuracy) metric is concealing this.

We calculated the correlation coefficient between the metrics and realized that GSMRE was better correlated with the actual MAD while MAPE was not.

	Correlation Coefficient
GSMRE, MAD	97%
MAPE, MAD	15%

Figure 32: Correlation coefficient between metrics for babywear.

Case Study Outcomes

The correlation coefficient of GSMRE/MAD was superior to MAPE/MAD. Using GSMRE as the accuracy metric would have a direct impact on the actual units. Thus, we would be more comfortable recommending GSMRE, instead of MAPE, to be used when a) improving the forecasting model and b) when a business manager in charge of the operation would take decisions based on the forecast model and its accuracy metric.

When these managers are considering how much money would be saved if they invest in improving the forecast accuracy by 1%, GSMRE would be a better and more correlated metric to base the answer on.

Conclusions & Recommendations

Some of the benefits for each measure can be summarized as follows:

1. MAD is of very good use when treating actual units
2. MAPE is very good when dealing with aggregated and stable data
3. MSE is good when handling small differences
4. GSMRE is very good at a granular level and has an excellent correlation with MAD, so it can also be used when dealing with actual units.

We also recommend calculating correlations between the actual units which might be crucial due to the link to direct actions, and the state of the system (for example supply chain or warehouse planning), to this metric.

1. The metric applied needs to agree with the problem at hand and the industry in which we work, for example working with an FMCG versus a general merchandise or non-food online business, versus a supermarket that sells groceries can be measured differently.

2. GSMRE is a very convenient metric in many cases, as shown in our example above.
3. One must apply a recurrent algorithm to avoid memory overflowing in an industrialization setting.

There is an additional benefit to GSMRE which is that it calculates the bias coefficient automatically. So, one has 2 metrics in one go.

The following table shows a general guideline for the choice of metric to evaluate forecast.

Product Type	GSMRE	MAPE	MAD	MSE
Slow Mover	✓ Correlates positively with MAD	✗ Hides small differences as numbers are close to zero	This should be a point of reference in all cases	✓
Volatile	✓ Very good with this kind of product	Hides volatility and averages out numbers		✗ Highlights outliers which are common in volatile products
Seasonal	✓ Depends on season but in general a good metric	✓ Depends on season but in general a good metric		✓ Depends on season but in general a good metric
Stable	✓	✓		✓
High Volume	✓	✓ Is very good as the volume approaches large numbers		✓

Figure 33: General guidelines for the choice of metrics to evaluate forecasts.

About the Authors

SERGE PLATA

Senior Principal, AI & Automation; London

Serge is a certified six-sigma black belt, and a chartered scientist in the UK. He was awarded a fellowship from the UK Institute of Mathematics and has more than 15 years' experience building, leading and consulting for data programs from FTSE-100 companies to SMEs and start-ups. He is currently the Chief Decision Scientist for Infosys Consulting EMEA.

HAIDONG (HENRY) SUN

Principal, AI & Automation; Chicago

Henry has extensive experience in data analytics and cloud computing. For the last 15 years, he has led consulting engagements in consumer goods, retail, manufacturing, hi-tech, financial services, defense and other industries. While focusing on digital transformation, he has led projects in demand forecast, logistics, new market entry strategy and M&A for clients from all continents except Antarctica.

MAGDA RACZKIEWICZ

Consultant AI&A; Warsaw

Magda is a Data Science Consultant for Infosys Consulting EMEA. She has over 3 years of experience consulting, designing, and building AI solutions as well as leading projects. She holds a master's degree in Computer Science and Econometrics with NLP specialization and scientific scholarship.

End Notes

1 - Adrian, M., Jaffri, A., Feinberg, D., "Market Guide: Graph Database Management Solutions," Gartner, 24 May 2021.

2 - McCreary, D., "Enterprise Knowledge Graphs: Graph Storytelling," Graph Storytelling - Enterprise Knowledge Graphs (dmccreary.github.io), accessed 2022/06/23;

McCreary, D. "Enterprise Knowledge Graph Trends for 2021," Medium, Enterprise Knowledge Graph Trends for 2021 by Dan McCreary, Medium, 2021/01/01, accessed 2022/06/06;

McCreary, D., "Intel and Dell Lead Next-Gen Enterprise Knowledge Graph Funding," Medium, 28 Feb 2021, accessed 2022/06/06;

McCreary, D., "Rules for Knowledge Graph Rules," Medium, 5 December 2020, accessed 2022/06/06.

3 - Ekstrand, M.D., Riedel, J.T., Konstan, J.A., "Collaborative Filtering Recommender Systems," Foundations and Trends in Human-Computer Interaction, Volume 4, Issue 2February 2011 pp 81–173https://doi.org/10.1561/110000000.

4 - Gilmer, Justin; Schoenholz, Samuel S.; Riley, Patrick F.; Vinyals, Oriol; Dahl, George E. (2017-07-17). "Neural Message Passing for Quantum Chemistry." International Conference on Machine Learning. PMLR: 1263–1272. arXiv:1704.01212.

5 - Ma, Tengfei; Ferber, Patrick; Huo, Siyu; Chen, Jie; Katz, Michael (2020). "Online planner selection with graph neural networks and adaptive scheduling." Proceedings of the AAAI Conference on Artificial Intelligence. 34: 5077--5084;

Osanlou, Kevin; Bursuc, Andrei; Guettier, Christophe; Cazenave, Tristan; Jacopin, Eric (2019). "Optimal Solving of Constrained Path-Planning Problems with Graph Convolutional Networks and Optimized Tree Search." 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS): 3519--3525. doi:10.1109/IROS40897.2019.8968113

6 - Gao, Y., Li, Y., Lin, Y., "Deep Learning on Knowledge Graph for Recommender System: A Survey," ResearchGate, Deep Learning on Knowledge Graph for Recommender System: A Survey | by Yang Gao | ResearchGate, March 2020, accessed 2022/06/08

7 - Kipf, Thomas N; Welling, Max (2016). "Semi-supervised classification with graph convolutional networks." International Conference on Learning Representations. 5 (1): 61–80. arXiv:1609.02907. doi:10.1109/TNN.2008.2005605. PMID 19068426. S2CID 206756462

8 - Derrow-Pinion, Austin; She, Jennifer; Wong, David; Lange, Oliver; Hester, Todd; Perez, Luis; Nunkesser, Marc; Lee, Seongjae; Xueying Guo; Wiltshire, Brett; Battaglia, Peter; Gupta, Vishal; Li, Ang; Xu, Zhongwen; Sanchez-Gonzalez, Alvaro; Li, Yujia; Velickovic, Petar (2021). "ETA Prediction with Graph Neural Networks in Google Maps." Arxiv: 3–4

9 - Hemsoth, N., "The Graph Neural Network behind your ETA," 31 August 2021, accessed 2022/06/08

10 - Linton, I., "How to increase Revenue for a Retail Bank," Chron, How to increase Revenue for a Retail Bank | by Ian Linton | Chron, accessed 2022/06/08

11- Blaschka, T., "How the world's largest Banks use Advanced Graph Analytics to fight Fraud," TigerGraph, How the world's largest Banks use Advanced Graph Analytics to fight Fraud | by Todd Blaschka | TigerGraph, 23 January 2020, accessed 2022/06/08

12 - Hunger, Michael, & Lyon, William, "Analyzing the Panama Papers with ne04j: Data Models, Queries and More, 08 April 2016, Analyzing the Panama Papers with Neo4j: Data Models, Queries & More, accessed 2022/06/20

13 - Yuhanna, N., Leganza, G., and Weber, D., "The Forrester Wave™: Graph Data Platforms, Q4, 2020, The 12 Providers that Matter Most And How They Stack Up," Forrester, 16 November, 2020.

14 - Kaynak, Okyay. "The golden age of Artificial Intelligence." 22 Sept 2021. link.springer.com/article/10.1007/s44163-021-00009-x

15 - Daniel Zhang, Nestor Maslej, Erik Brynjolfsson, John Etchemendy, Terah Lyons, James Manyika, Helen Ngo, Juan Carlos Niebles, Michael Sellitto, Ellie Sakhaee, Yoav Shoham, Jack Clark, and Raymond Perrault, "The AI Index 2022 Annual Report," AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University, March 2022.

16 - Ibid.

17 - Ibid.

18 - Gartner. "Gartner Identifies Four Trends Driving Near-Term Artificial Intelligence Innovation. 7 Sept 2021. www.gartner.com/en/newsroom/press-releases/2021-09-07-gartner-identifies-four-trends-driving-near-term-artificial-intelligence-innovation

19 - Office of the Chief Economist. "Inventing AI." US Patent and Trademark Office. Oct 2020. www.uspto.gov/sites/default/files/documents/OCE-DH-AI.pdfo.gov

20 - Will Chen, Pengju Shang, Yunlai Zha. "Patent Protection on AI Inventions." 30 Aug 2021. www.intellectualpropertylawblog.com/archives/patent-protection-on-ai-inventions

21 - Bernard, Allen. TechRepublic. "59% of Senior Executives Feel Threatened by Artificial Intelligence." 4 June 2021.

22 - Sam Ransbotham, Shervin Khodabandeh, Ronny Fehling, Burt Lafountain, David Kiron. "Winning with AI." MIT Sloan Management Review. 15 Oct 2019

23 - "Data and AI Leadership Executive Survey." NewVantage Partners. 2022. www.newvantage.com/_files/ugd/e5361a_2f859f3457f24cff9b2f8a2bf54f82b7.pdf

24 - Bernard, Allen. TechRepublic. "59% of Senior Executives Feel Threatened by Artificial Intelligence." 4 June 2021.

25 - Migdal, Miki. "The Vertical Approach: Why AI Should Be Industry Specific," Forbes. 28 May 2021. www.forbes.com/sites/forbestechcouncil/2021/05/28/the-vertical-approach-why-ai-should-be-industry-specific/?sh=3f6b4cb320ce

26 - Rizzoli, Alberto. "8 Practical Applications of AI in Agriculture. V7, 6 June 2022. www.v7labs.com/blog/ai-in-agriculture

27 - Imperial College London. "Low-cost AI soil sensors could help farmers curb fertilizer use." 13 Dec 2021. www.sciencedaily.com/releases/2021/12/211213111704.htm

28 - Ibid.

29 - OECD. "Managing water sustainably is key to the future of food and agriculture." www.oecd.org/agriculture/topics/water-and-agriculture

30 - "Farmers in India are Using AI to Increase Crop Yields." India Agro Blog. www.indiaagronet.com/Agriculture-Technology/Farmers-in-India-are-Using-AI-to-Increase-Crop-Yields.html

31 - "Data and AI Leadership Executive Survey." NewVantage Partners. 2022. www.newvantage.com/_files/ugd/e5361a_2f859f3457f24cff9b2f8a2bf54f82b7.pdf

32 - Bernard, Allen. TechRepublic. "59% of Senior Executives Feel Threatened by Artificial Intelligence." 4 June 2021.

End Notes

33 - PwC, "2022 AI Business Survey." www.pwc.com/us/en/tech-effect/ai-analytics/ai-business-survey.html

34 - Zaman, Kashif. Forbes. "How to Offer New Human+Machine Value Propositions with AI." 11 Jun 2020. www.forbes.com/sites/forbesagencycouncil/2020/06/11/how-to-offer-new-humanmachine-value-propositions-with-ai/?sh=5170b029240c

35 - Bernard, Allen. TechRepublic. "59% of Senior Executives Feel Threatened by Artificial Intelligence." 4 June 2021.

36 - Daniel Zhang, Nestor Maslej, Erik Brynjolfsson, John Etchemendy, Terah Lyons, James Manyika, Helen Ngo, Juan Carlos Niebles, Michael Sellitto, Ellie Sakhaee, Yoav Shoham, Jack Clark, and Raymond Perrault, "The AI Index 2022 Annual Report," AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University, March 2022.

37 - Ibid.

38 - Infosys Knowledge Institute. www.infosys.com/content/dam/infosys-web/en/techcompass/responsible-ai.html

39 - Hall, T., & Yoon, E. (2017, August 29). Searching for New Ideas in the Curious Things Your Customers Do. Harvard Business Review. Retrieved June 27, 2022, from https://hbr.org/2017/04/searching-for-new-ideas-in-the-curious-things-your-customers-do

40 - Alford, J. (2021, August 6). GDPR and Artificial Intelligence. The Regulatory Review. Retrieved June 27, 2022, from https://www.theregreview.org/2020/05/09/saturday-seminar-gdpr-artificial-intelligence/

41 - Lopez de Prado Marcos. The 10 reasons most machine learning funds fail, The Journal of Portfolio Management (2018), 4496) 120-133

42 - Agrawal, A., Gans, J., & Goldfarb, A. (2020, August 18). How to Win with Machine Learning. Harvard Business Review. Retrieved June 27, 2022, from https://hbr.org/2020/09/how-to-win-with-machine-learning; Ilik, S. (2022, January 7). Machine Learning and the GDPR - DataBulls. Medium. Retrieved June 27, 2022, from https://medium.com/databulls/machine-learning-and-the-gdpr-253a627e569f

43 - Babic, B., Cohen, I., Evgeniou, T., & Gerke, S. (2020, December 15). When Machine Learning Goes Off the Rails. Harvard Business Review. Retrieved June 27, 2022, from https://hbr.org/2021/01/when-machine-learning-goes-off-the-rails

44 - Ved, A. (2021, November 15). How to develop Artificial Intelligence that is GDPR-friendly. TechGDPR. Retrieved June 27, 2022, from https://techgdpr.com/blog/develop-artificial-intelligence-ai-gdpr-friendly/

45 - Arnold, M. J., & Reynolds, K. E. (2003). Hedonic shopping motivations. Journal of retailing, 79(2), 77-95

46 - Shanes, Justin on. (2016, November 29). Twitter. Retrieved June 27, 2022, from https://twitter.com/justinshanes/status/803453049603690496

47 - Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. Advances in neural information processing systems, 26

48 - VanderPlas, J. (2017). Python Data Science Handbook. O'Reilly Media Inc, Sebastopol Ca. Retrieved June 27, 2022, from https://books.google.nl/books?id=6omNDQAAQBAJ&redir_esc=y

49 - Sartor, G. & European Parliament. European Parliamentary Research Service. Scientific Foresight Unit. (2020). The Impact of the General Data Protection Regulation (GDPR) on Artificial Intelligence. European Parliament

50 - The definition of the GSMRE is where 'e' is the error defined as is a better metric than the standard MAPE.

51 - The formal for MAPE is $\text{MAPE} = \frac{1}{n}\sum \left| \frac{actual - forecast}{actual} \right|$

52 - Lawrence Michael, O'Connor Marcus, Judgemental forecasting in the presence of loss functions, International Journal of Forecasting 21, 3-14 (2005)

53 - See https://peerj.com/preprints/3190/ which can also be used.

Additional References

Doubling Down on Scaling Your AI

Davenport, Tom. "Return on Artificial Intelligence: The Challenge and the Opportunity." Forbes. 27 Mar 2020. www.forbes.com/sites/tomdavenport/2020/03/27/return-on-artificial-intelligence-the-challenge-and-the-opportunity/?sh=3c82d3386f7c

Goasduff, Laurence. "The 4 Trends That Prevail on the Gartner Hype Cycle for AI, 2021." Gartner. 22 Sept 2021. www.gartner.com/en/articles/the-4-trends-that-prevail-on-the-gartner-hype-cycle-for-ai-2021

Reese, Hope. "AI Remains Priority for CEOs, According to New Gartner Survey." 19 May, 2022. TechRepublic. www.techrepublic.com/article/ai-remains-priority-ceos-gartner-survey/

Bernard, Allen. "59% of senior executives feel threatened by artificial intelligence." 4 Jun 2021. TechRepublic. www.techrepublic.com/article/ai-gaining-traction-in-the-workplace

Vishal Dhupar, Rafee Tarafdar, Dr. Shailesh Kumar. "Designing & Deploying an enterprise AI Strategy." Webcast, ET CIO AI Leaders Forum. 25 Jan 2021. www.youtube.com/watch?v=ZaR7SHKniP8

Jacomo Corbo, David Harvey, Nicolas Hohn, Kia Javanmardian, Nayur Khan. "Scaling AI like a tech native: The CEO's role." 13 Oct 2021. www.mckinsey.com/business-functions/quantumblack/our-insights/scaling-ai-like-a-tech-native-the-ceos-role

Infosys TechCompass. Infosys Knowledge Institute. www.infosys.com/content/dam/infosys-web/en/techcompass/computer-vision.html

Corey McNair, Sean Owens, Margo Visitacion. "Effectiveness Is the New Metric for Tech Value." 23 Jun 2021. Forrester. www.forrester.com/report/Effectiveness-Is-The-New-Metric-For-Tech-Value/RES162936

McKendrick, Joe. "AI Adoption Skyrocketed Over the Past 18 Months." 27 Sep 2021. Harvard Business Review. www.hbr.org/2021/09/ai-adoption-skyrocketed-over-the-last-18-months

Thomas H. Davenport and Randy Bean. "Companies Are Making Serious Money with AI." 17 Feb 2022. MIT Sloan Management Review. www.sloanreview.mit.edu/article/companies-are-making-serious-money-with-ai/

Anie Juby, Ashok Panda. "AI For the Rest of the Enterprise." www.infosys.com/insights/ai-automation/ai-rest-enterprise.html

Armstrong, J. S., Collopy, F. (1992). Error measures for generalizing about forecasting methods: Empirical comparisons. International Journal of Forecasting, 8, 69 – 80.

Armstrong, J. S. (2001). Evaluating forecasting methods, In J. S. Armstrong (Ed.), Principles of forecasting: A handbook for researchers and practitioners. Norwell, Kluwer Academic Publishers.

Batseiler Jordy, Vanhoucke Mario, 2017, Improving project forecast accuracy by integrating earned value management with exponential smoothing and reference class forecasting, International Journal of Project management 35 (2017) 28-43

Box, G.E.P. and Jenkins, G.M. 1976. Time Series Analysis: Forecasting and Control, Revised Edn., Hoden-Day, San Francisco.

Box, G.E.P., Jenkins, G.M. and Reinsel, G.C. 1994. Time Series Analysis, Forecasting and Control, Englewood Cliffs, N.J: Prentice Hall.

Chatfield, C. (1988). Apples, oranges and mean square error. International Journal of Forecasting, 4, 515 – 518.

Davydenko, A., Fildes, R., 2013. Measuring forecasting accuracy: The case of judgmental adjustments to SKU-level demand forecasts. International Journal of Forecasting 29 (3), 510–522.

Deyslen Mariano-Hernández, Luis Hernández-Callejo, Felix Santos García, Oscar Duque-Perez, Angel L. Zorita-Lamadrid, A Review of Energy Consumption Forecasting in Smart Buildings: Methods, Input Variables, Forecasting Horizon and Metrics, Applied Science. 2020, 10, 8323, available at www.mdpi.com/journal/applsci

Elliot G. Timmermann A. 2008, Economic forecasting, Journal of Economic Literature 46 (1) 3-56

Elshaer, R., 2013. Impact of sensitivity information on the prediction of project's duration using earned schedule method. Int. J. Proj. Manag. 31, 579–588.

Fildes, R., Goodwin, P., Lawrence, M., Nikolopoulos, K., 2009. Effective forecasting and judgmental adjustments: an empirical evaluation and strategies for improvement in supply-chain planning. International Journal of Forecasting 25, 3–23.

Holt, C., 1957. Forecasting Seasonal and Trends by exponentially Weighted Moving Averages. Office of Naval Research Memorandum.

Holt, C., 2004. Forecasting Seasonal and Trends by Exponentially Weighted Moving Averages. Int. J. Forecast. 20, 5–10.

Holt, C., Modigliani, F., Muth, J., Simon, H., 1960. Planning Production, Inventories, and Work Force. Prentice-Hall, Englewood Cliffs, New Jersey.

Hull John C., 2005, Fundamentals of Futures and Options Markets, Pearson Education Ltd, London

Hyndman Rob J., Koehler Anne B., Another look at measures of forecast accuracy, international journal of forecasting 22 (2006) 679-688

Khamooshi, H., Golafshani, H., 2014. EDM: earned duration management, a new approach to schedule performance management and measurement. Int. J. Proj. Manag. 32, 1019–1041.

Kourentzes Nikolaos, Trapero Juan R. Svetunkov Ivan., 2014, Measuring the behavior of experts on demand forecasting: A complex task, working paper, Lancaster University Management School, Lancaster University, 2014

Kremer, M., Moritz, B., Siemsen, E., 2011.Demand forecasting behavior: System neglect and change detection. Management Science 57 (10), 1827– 1843.

Lawerence Michael, O'Connor Marcus, 2005, Judgmental forecasting in the presence of loss functions, International Journal of Forecasting 21, (2005) 3-14

Lipke, W., 2011. Schedule adherence and rework. PM World Today 13, 1–14. 9

Makridakis, S., Anderson, A., Carbone, R., Fildes, R., Hibon, M., Lewandowski, R., et al. (1982). The accuracy of extrapolation (time series) methods: Results of a forecasting competition. Journal of Forecasting, 1, 111 – 153.

McKenzie Jordi, Mean absolute percentage error and bias in economic forecasting, Economic Letters 113 (2011) 259-262

Trapero, J. R., Fildes, R., Davydenko, A., 2011. Nonlinear identification of judgmental forecasts effects at SKU level. Journal of Forecasting 30, 490–508.

Vanwinckelen Gitte, Blockeel Hendrik, 2012, Journal BeneLearn 2012: Proceedings of the 21st Belgian-Dutch conference on machine learning, 39-44
Winters, P., 1960. Forecasting sales by exponentially weighted moving averages. Manag. Sci. 6, 324–342.

Wheatcroft Edward, Interpreting the skill score form of forecast performance metrics, International Journal of Forecasting 35 (2019) 573–579

Infosys® | CONSULTING

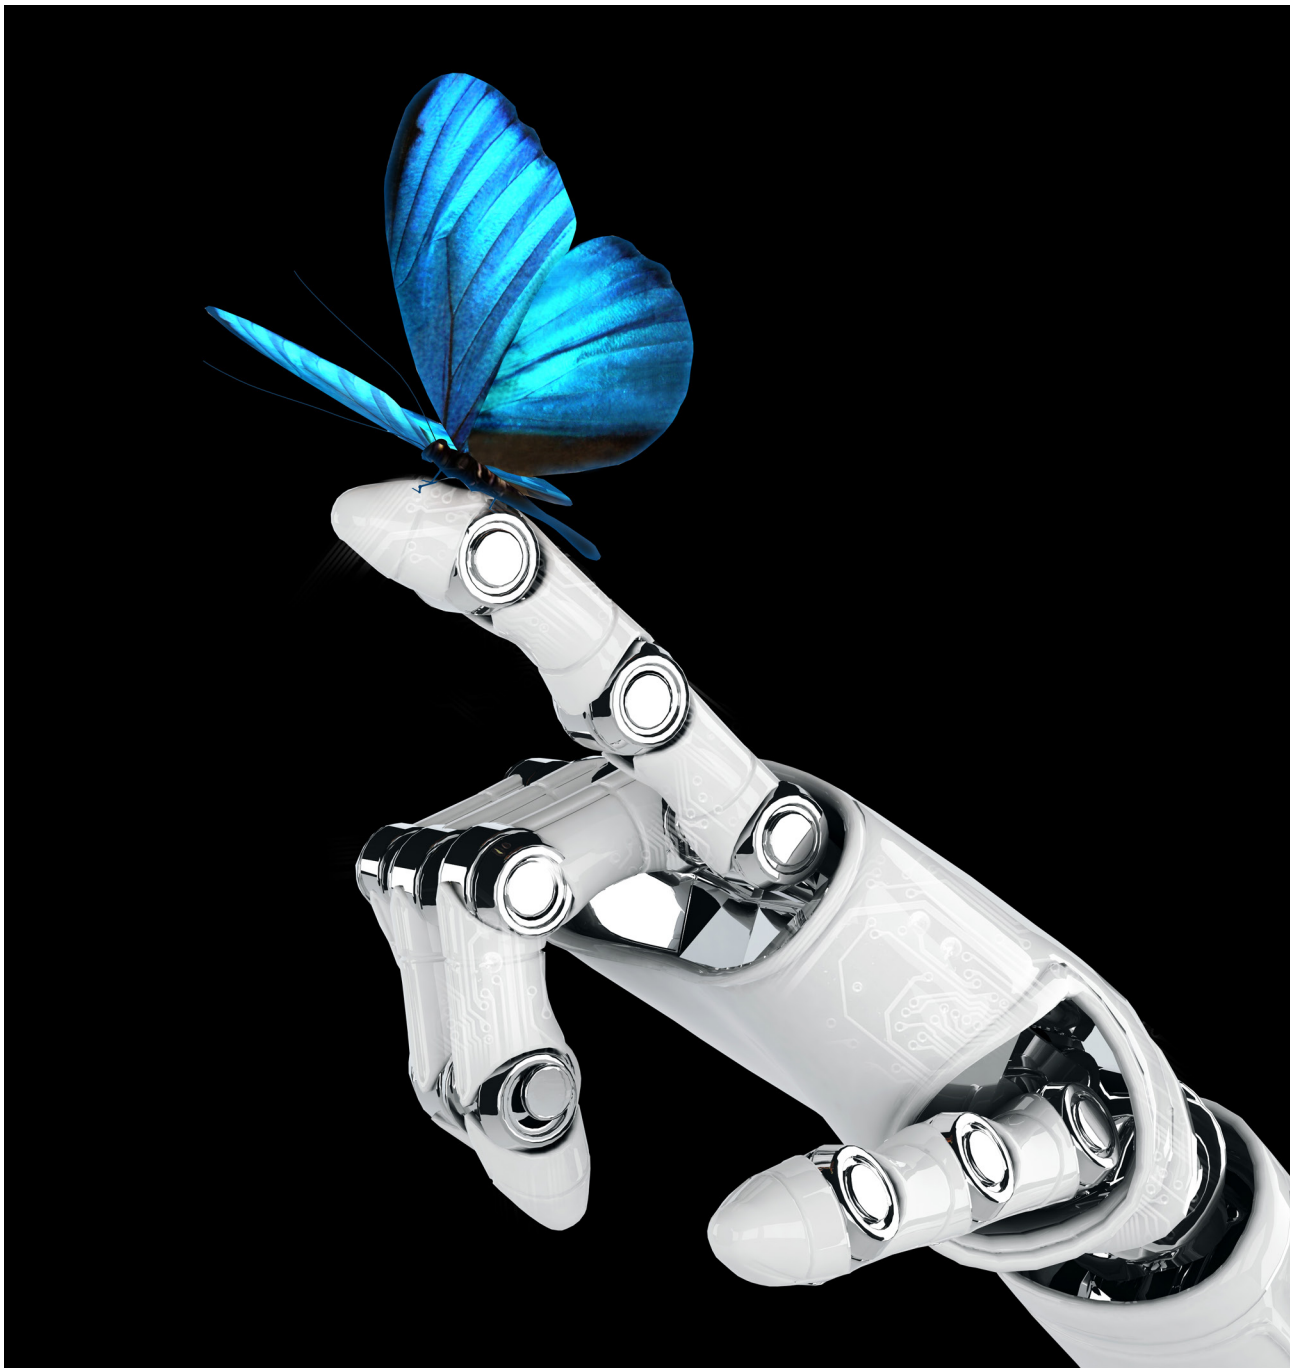

SETH SIEGEL

MANAGING PARTNER, AI, ANALYTICS & CYBER SECURITY

seth.siegel@infosys.com

DAVID SEMACH

EMEA HEAD OF AI & AUTOMATION

david_semach@infosys.com